aruba®

NETWORKS

ArubaOS-CX Switching Fundamentals, rev 20.22

LAB GUIDE

EDU-CXF-RLABS-v20.22

ArubaOS-CX Switching Fundamentals

20.22

Lab Guide

April 2022

ArubaOS-CX Switching Fundamentals, rev 20.22

Legal Notice
The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty
This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.
Altering this device (such as painting it) voids the warranty.

SKU: EDU-CXF-RLABS-20.22

ArubaOS-CX Switching Fundamentals Lab Guide

Table of Contents

AOS-CX Switching Fundamentals

Lab 0: Testing Lab Connectivity (optional)

Overview

The Aruba Training Lab provides you with the equipment you need for completing several lab activities. You should know the purpose and access procedures to this equipment.

- **PC-1**: This client is used for traffic analysis, connectivity testing, accessing the Web-UI of NetEdit and your switches and accessing the CLI over SSH of Core switches, ISP1 and Server Switch.
- **PC-3**: This client is used connectivity testing.
- **PC-4**: This client is used connectivity testing.
- **6300-A switch**: This is one of your Access-Switches.
- **6300-B switch**: This is one of your Access-Switches.
- **Core-1 switch**: This is a shared resource and you will access it over an SSH session via PC-1.
- **Core-2 switch**: This is a shared resource and you will access it over an SSH session via PC-1.
- **OOBM switch**: You have NO access to this switch.
- **ISP1**: This is a shared resource and you will access it over an SSH session via PC-1.
- **ISP2**: You have NO access to this switch.
- **Server Switch**. This is a shared resource and you will access it over an SSH session via PC-1.
- **NetEdit**: You will access this device over a HTTPS session via PC-1
- **Windows Server**: You have NO access to this server, but you will access its web page and download files running TFTP from PC-1.
- **ClearPass server**: You have NO access to this sever, but you will use it as a AAA server for your switches.

Objectives

After completing this lab, you will have all the information needed to support the hands-on labs in this course.

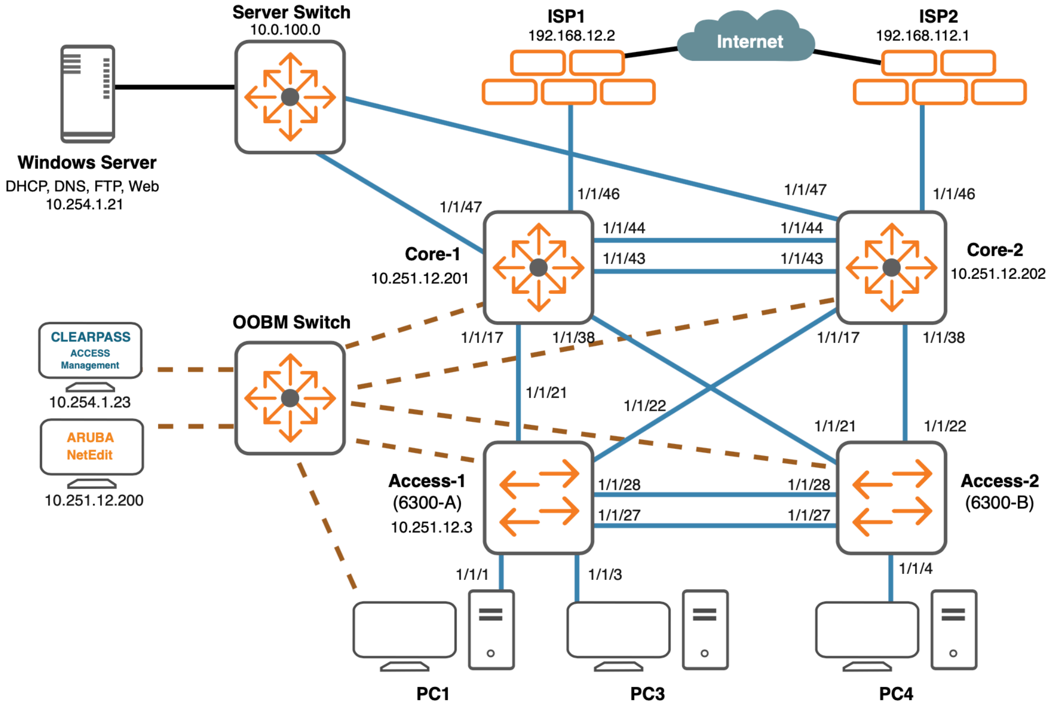

Figure 0-1: Lab Topology

Task 1: Aruba Training Lab Access

Objectives

To check that you have connectivity to the remote lab and can successfully login. This will ensure that you have access to your remote lab equipment during this training.

Steps

1. On your local computer, launch a web browser, and enter to the Aruba Training Lab web portal at the URL:https://arubatraininglab.computerdata.com.
2. Enter your **username** and **password** (if you do not have one, ask your instructor for the credentials), and click the **Sign in** button.

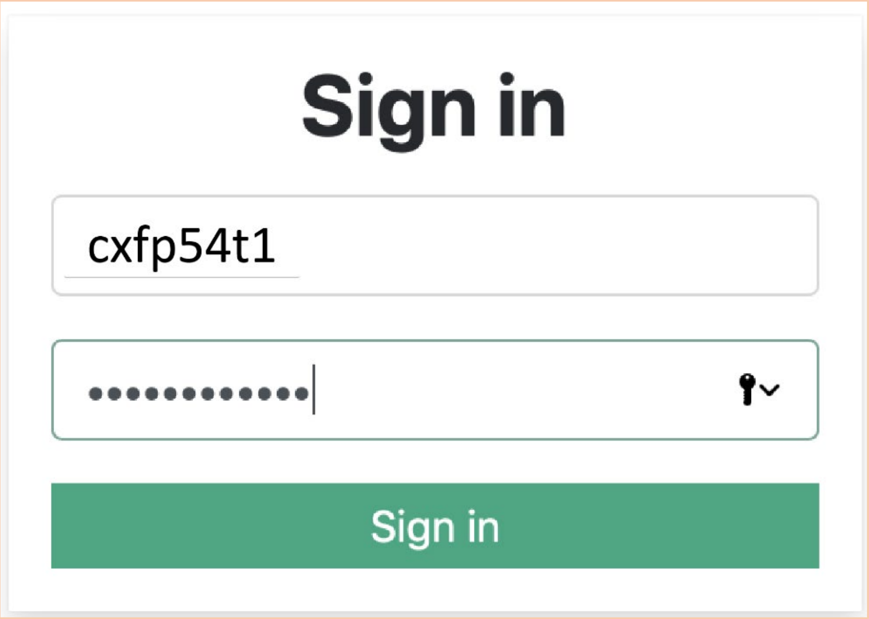

Sign in

cxfp54t1

••••••••••••

Figure 0-2: Sign in

Task 2: Testing Connectivity

Objectives

To test connectivity and authentication credentials for each of the devices. Working from the Aruba Training Lab diagram, you will connect to and log into the Access switches and your client PCs.

6300-A and 6300-B

1. To connect to the console of the 6300-A switch, right-click on the icon in the lab diagram and select "**Open Console.**"

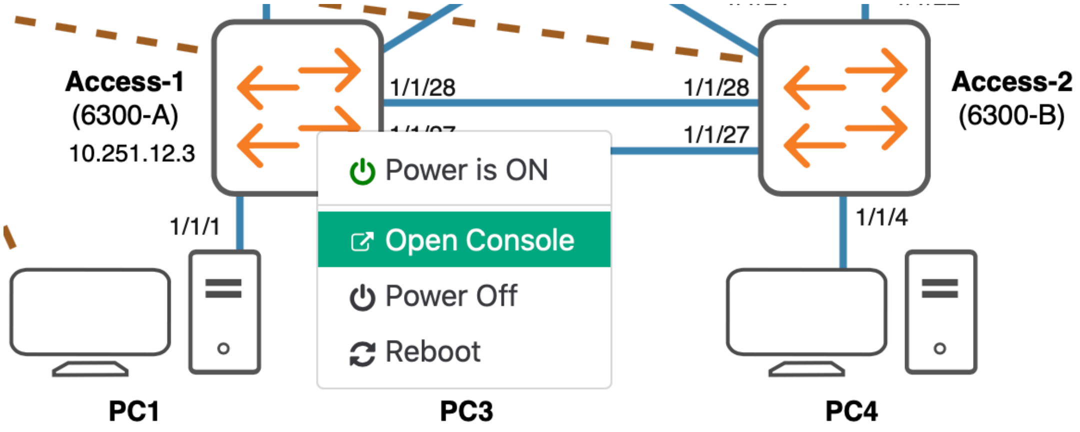

Figure 0-3: Open Console of 6300-A

2. A new browser tab should open with a blank, black screen.
3. Press **[enter]** a couple times, and you will see a user prompt.
4. Login using **admin** and **no password**.
5. It will ask you to define a new password, hit **[enter]** twice.

```
6300 login: admin
Password:


Please configure the 'admin' user account password.
Enter new password:
Confirm new password:
```

```
6300#
```

6. Repeat steps 1 to 5 on 6300-B.

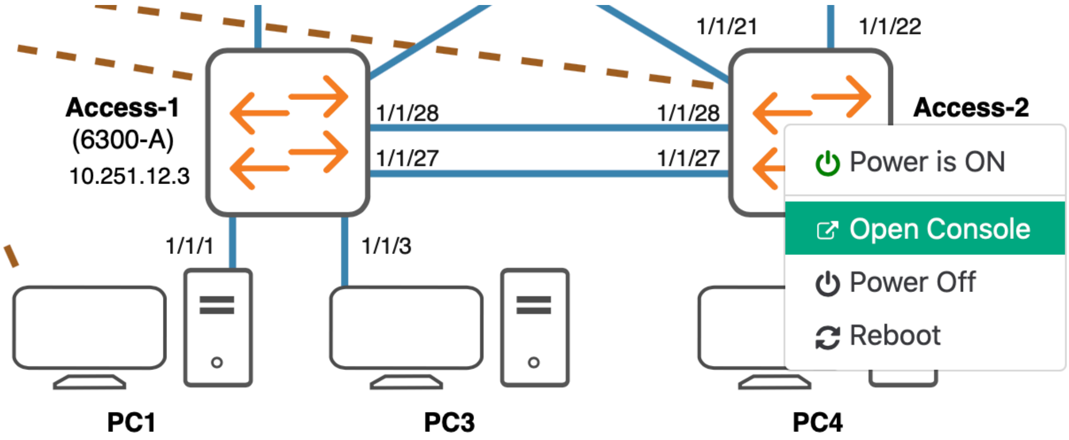

Figure 0-4: Open Console of 6300-B

```
6300 login: admin
Password:


Please configure the 'admin' user account password.
Enter new password:
Confirm new password:
6300#
```

PC-1, PC-3 and PC-4

7. To access the desktop PC-1, just Right-click on the icon in the lab diagram and select "**Open Desktop**."

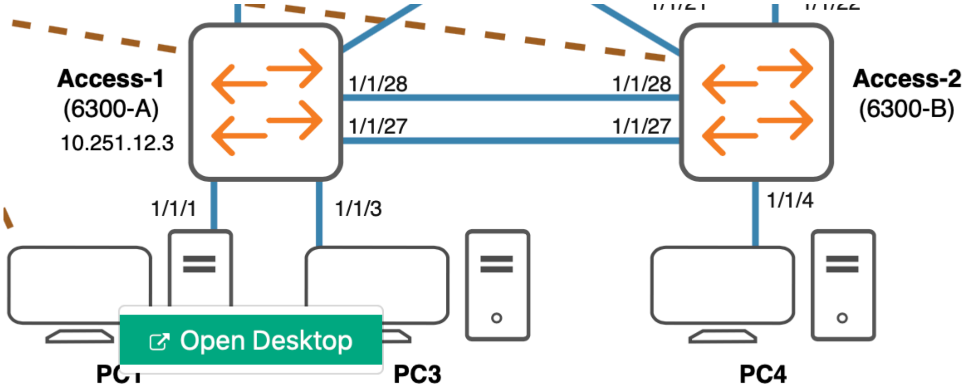

Figure 0-5: Open Console of PC-1

8. A new browser tab will open with the remote desktop.

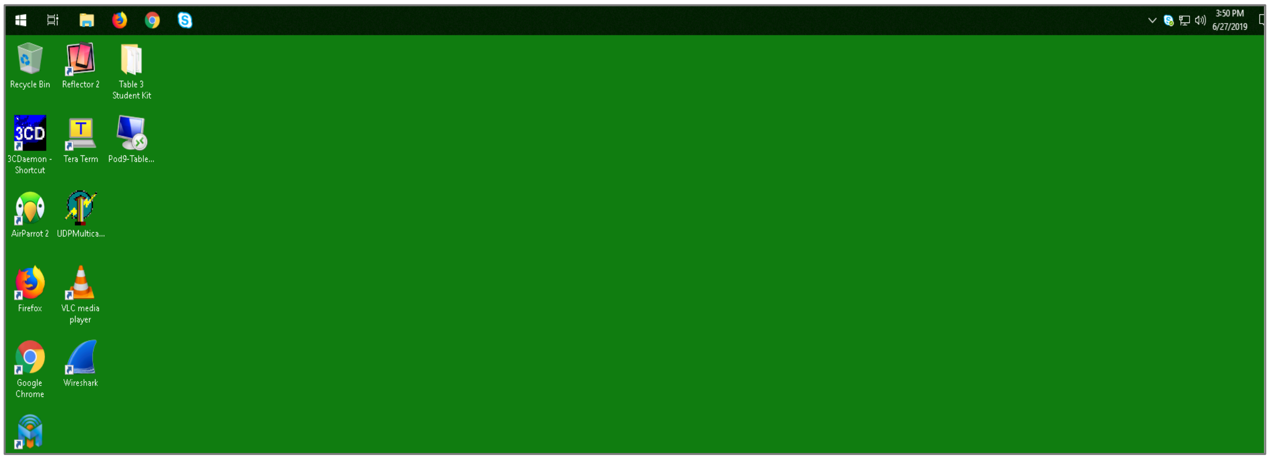

Figure 0-6: PC-1’s desktop

NOTE: It may take a few minutes for the PC-1 desktop to come up. Also, if your Aruba Training Lab has been idle for a while after you login, you may need to log out of the lab interface and log back in and then launch the desktop again.

9. Repeat steps 7 and 8 on PC-3 and PC-4.

Core-1 (via PC-1)

10. Move back to PC-1.

11. Open Putty. You will find saved sessions to Core-1 and other three devices.

TIP: Putty should have Saved Sessions to Core-1 and Core-2, you could use those as a shortcut.

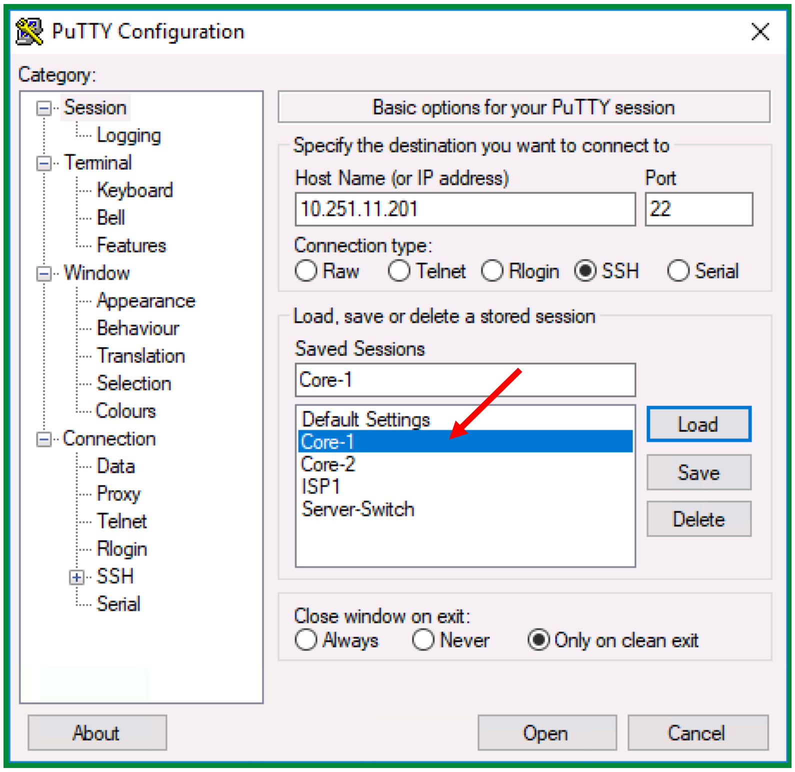

Specify the destination you want to connect to

Host Name (or IP address): 10.251.11.201 Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Load, save or delete a stored session

Saved Sessions: Core-1

Default Settings
Core-1
Core-2
ISP1
Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 0-7: PC-1's desktop

12. Double click **Core-1** saved session.
13. Login using **cxfX/aruba123**

NOTE: Replace the highlighted "X" with your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

Core-2, ISP1 and Server Switch (via PC-1)

14. Repeat steps 11 to 13 for Core-2, ISP1 and Server-Switch.

NetEdit (via PC-1)

15. Move back to **PC-1**.
16. Open a browser and type the NetEdit IP address in the URL field (**10.251.X.200**) then hit **[enter]**. You will be presented a security certificate warning.
17. Accept the warning. You will see the NetEdit login page right after.

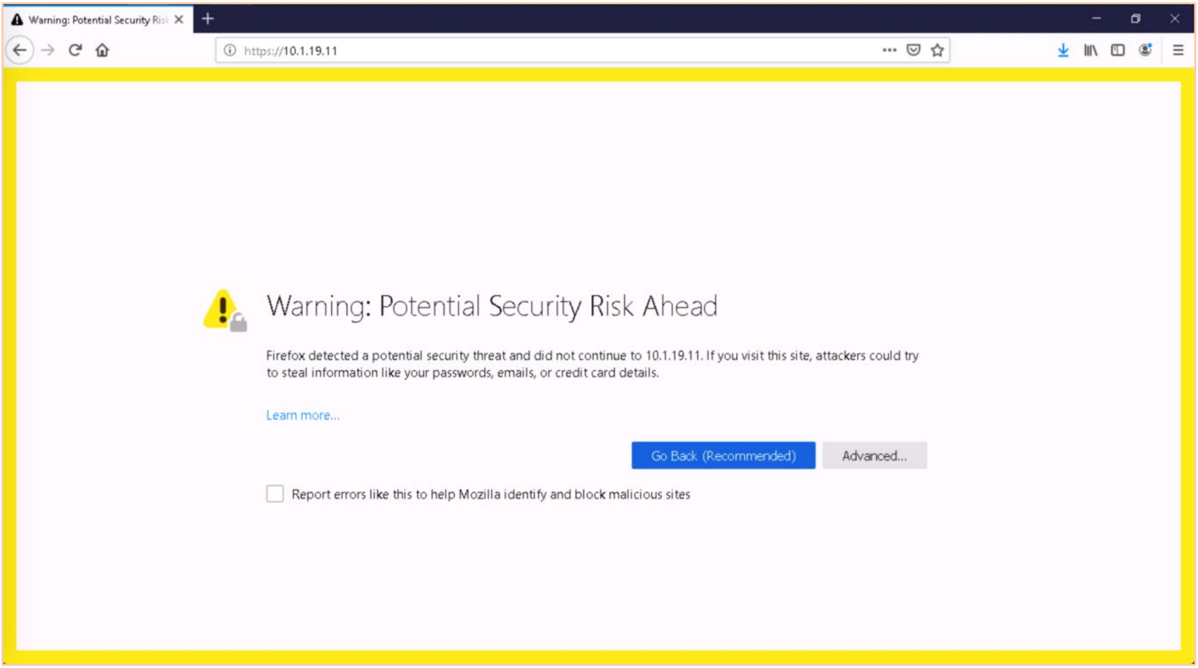

Warning: Potential Security Risk Ahead

Firefox detected a potential security threat and did not continue to 10.1.19.11. If you visit this site, attackers could try to steal information like your passwords, emails, or credit card details.

□ Report errors like this to help Mozilla identify and block malicious sites

Figure 0-8: Digital Certificate Warning

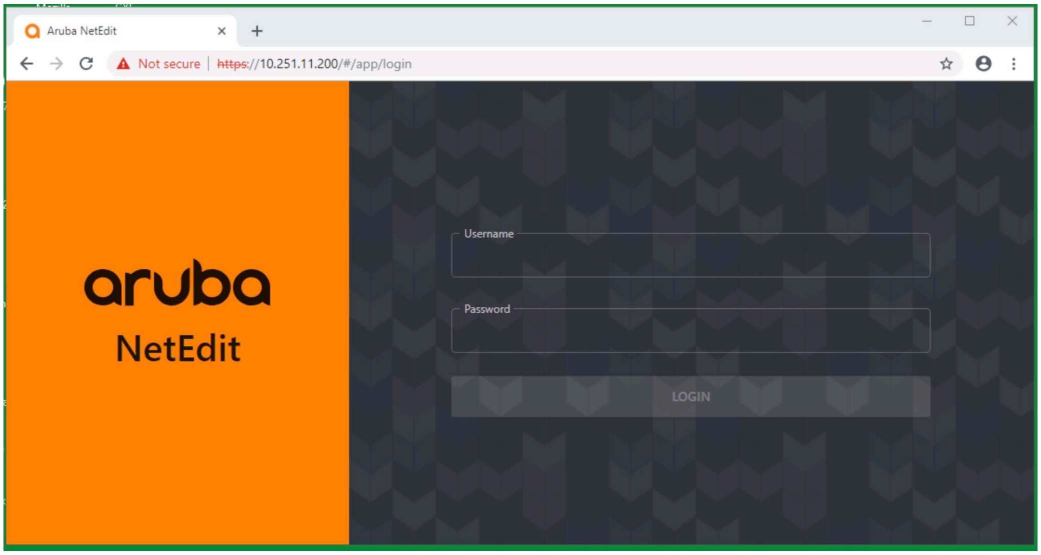

Figure 0-9: NetEdit Login Page

NOTE: If accessing any device is not successful then notify your instructor.

You have completed Lab 0!

AOS-CX Switching Fundamentals

Lab 1: Numerical Conversion

Overview

Welcome to the Aruba Switching Fundamentals training course. This lab manual will be your companion in your networking education journey. It contains different activities such as configuration, debugging and verification, troubleshooting, topology discovery, subnetting, traffic analysis, demonstrations and more, with the main goal of sharing with you the knowledge and required skills for deploying a medium sized single site campus network using AOS-CX switching platforms.

Although this training assumes no previous networking knowledge and is intended to transmit solid fundamental concepts, some tasks will cover details in depth, from the ground up.

With the exception of Lab 1, the rest of the book will take you into a scenario where a company called BigStartup needs your professional networking services to achieve business success.

Be aware that you may have to work after hours to complete all of the labs.

The current lab is limited to practicing some binary and hexadecimal conversions.

Objectives

After completing this lab, you will be able to:

- Convert decimal numbers to binary, hexadecimal and vice versa

Task 1: Binary to Decimal Conversion

Objectives

Convert the following binary into decimal values.

a) 10101010
b) 11100011
c) 01110000 (optional)
d) 10000001 (optional)
e) 00011100 (optional)

Steps

1. Fill out Table 1-1 with the “Power of two” information shown in Module 1 – Numerical Systems.
2. Use table 1-1 for completing your conversions.

TIP: In your time off, practice writing the table down. The more times you do it the easier for you to remember it. This is a good shortcut for Decimal to binary conversion whenever a calculator isn’t close.

Table 1-1: Power of 2: Binary to decimal

Powers of 2								
Decimal								

Binary a)								
Decimal a)								

Binary b)								
Decimal b)								

Binary c)								
Decimal c)								

Binary d)								
Decimal d)								

Binary e)								
Decimal e)								

Task 2: Decimal to Binary Conversion Method 1

Objectives

Convert the following decimal values into binary using the division method:

a) 315
b) 116
c) 39 (optional)
d) 240 (optional)

Steps

1. Convert 315

Rem. 9	Rem. 8	Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1

2. Convert 116

Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1

3. Convert 39

Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1

4. Convert 240

Rem. 8	Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1

Task 3: Decimal to Binary Conversion Method 2

Objectives

Convert the following decimal values into binary using the power of two method.

a) 224
b) 17
c) 199 (optional)
d) 46 (optional)

Steps

1. Fill out Table 1-2 with the “power of two” information shown in Module 1 – Numerical Systems.
2. Use table 1-2 for completing your conversions.

Table 1-2: Power of 2: Decimal to binary

Power of 2								
Decimal								
Binary a)								
Binary b)								
Binary c)								
Binary d)								

Task 4: Binary to Hexadecimal Conversion

Objectives

Convert the following binary values into hexadecimal.

a) 01100110
b) 10100101
c) 00010010 (optional)
d) 01011010 (optional)

Steps

1. Fill out Table 1-3 with the “Decimal to Hexadecimal” information shown in Module 1 – Numerical Systems.
2. Use table 1-3 for completing your conversions.

Table 1-3: Binary to Hexadecimal

Binary	Hexadecimal

3. Convert 01100110

4. Convert 10100101

5. Convert 00010010

6. Convert 01011010

Task 5: Hexadecimal to Binary Conversion

Objectives

Convert the following hexadecimal values into binary using the division method.

a) AB
b) AB3
c) 3F4 (optional)
d) 0C (optional)

Steps

1. Convert AB

2. Convert AB3

3. Convert 3F4

4. Convert 0C

Task 6: Decimal to Hexadecimal Conversion (optional)

Objectives

Convert the following decimal values into hexadecimal using the division method.

a) 898
b) 2033
c) 1572
d) 78

Steps

1. Fill out Table 1-4 with the “Decimal to Hexadecimal” information shown in Module 1 – Numerical Systems.
2. Use table 1-4 for completing your conversions.

Table 1-4: Decimal to Hexadecimal

Decimal	Hexadecimal

3. Convert 898

Rem. 3	Rem. 2	Rem. 1

4. Convert 2033

Rem. 3	Rem. 2	Rem. 1

5. Convert 1572

Rem. 3	Rem. 2	Rem. 1

6. Convert 78

Rem. 2	Rem. 1

Task 7: Hexadecimal to Decimal Conversion (optional)

Objectives

Convert the following hexadecimal values into binary using the division method.

a) F3A
b) 15B
c) 111
d) 7C

Steps

1. Fill out Table 1-5 with the “Power of 16” information shown in Module 1 – Numerical Systems.
2. Use the table 1-4 and 1-5 for completing your conversions.

Table 1-5: Decimal to Hexadecimal

Powers of 16				
Decimal				

Hexadecimal a)				
Decimal				
Multiplication				
Decimal a)				

Hexadecimal b)				
Decimal				
Multiplication				
Decimal b)				

Hexadecimal c)				
Decimal				
Multiplication				
Decimal c)				

Hexadecimal d)				
Decimal				
Multiplication				
Decimal d)				

You have completed Lab 1!

AOS-CX Switching Fundamentals

Lab 2: Packet Exploration

Overview

In the current lab you will explore Ethernet, IP, TCP and UDP packet headers and be familiar with their contents.

Objectives

After completing this lab, you will be able to:

- Capture packets using Wireshark
- Explore layer 2, 3 and 4 headers
- Identify most significant fields in headers

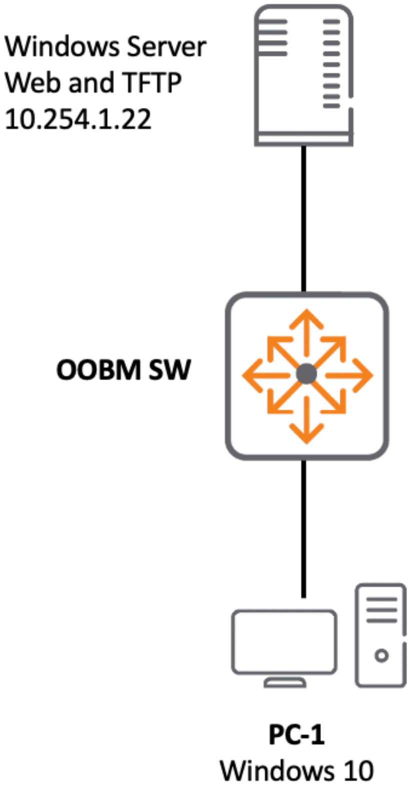

Figure 2-1: Lab Topology

Task 1: Discover Headers and Encapsulation

Objectives

A key step for better learning data forwarding and networking protocols is being able to look at packets and identify their OSI model headers, and the headers contents.

In this task you will explore Ethernet, IP, UDP and TCP headers.

Steps

PC-1

1. Open a console session to **PC-1**.
2. Open **Wireshark**, there should be a shortcut on the Desktop.

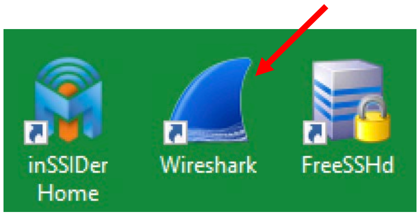

Figure 2-2: Wireshark shortcut

NOTE: Wireshark is a well-known, open source packet analyzer tool. It is capable of capturing traffic in different media types such as Ethernet, 802.11, Bluetooth, USB and more. It is supported on main desktop operating systems such as Microsoft Windows, MacOS and many Linux distributions. For more information please go to:

www.wireshark.org
https://wikipedia.org/wiki/Wireshark

3. Expand the "**View**" menu and uncheck the "**Packet Bytes**" option.

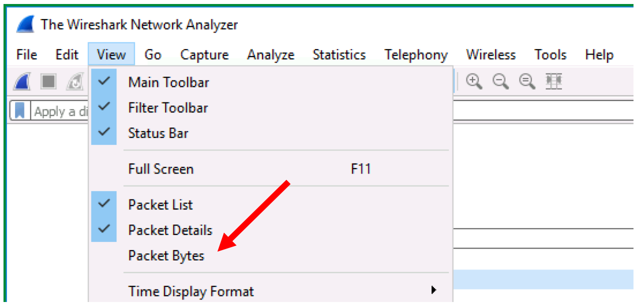

Figure 2-3: Wireshark View > Packet Bytes

4. Double click the **OOBM** entry. That will begin the packet capture in that interface.

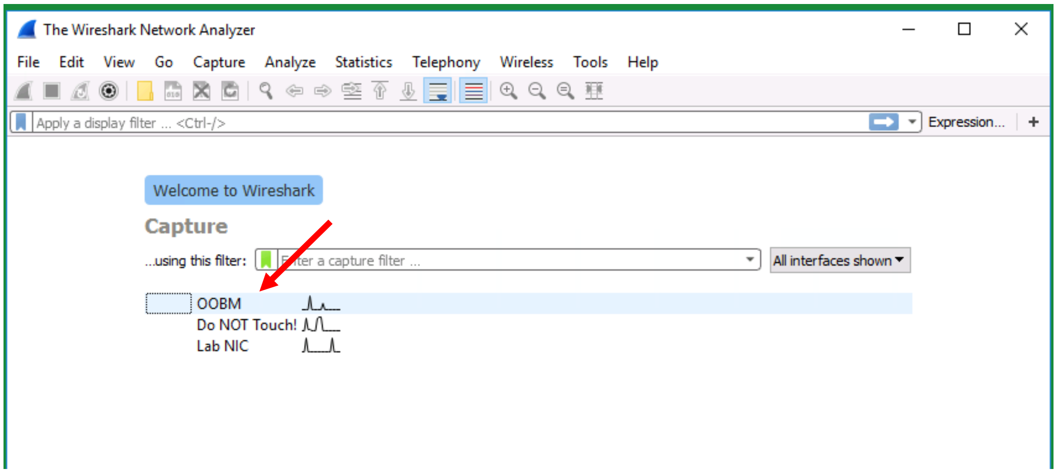
The Wireshark Network Analyzer
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Apply a display filter ... <Ctrl-/>
Expression...
Welcome to Wireshark
Capture
...using this filter:
Enter a capture filter ...
All interfaces shown

OOBM
Do NOT Touch!
Lab NIC

Figure 2-4: Wireshark NICs

5. Identify the components shown in figure 2-5.

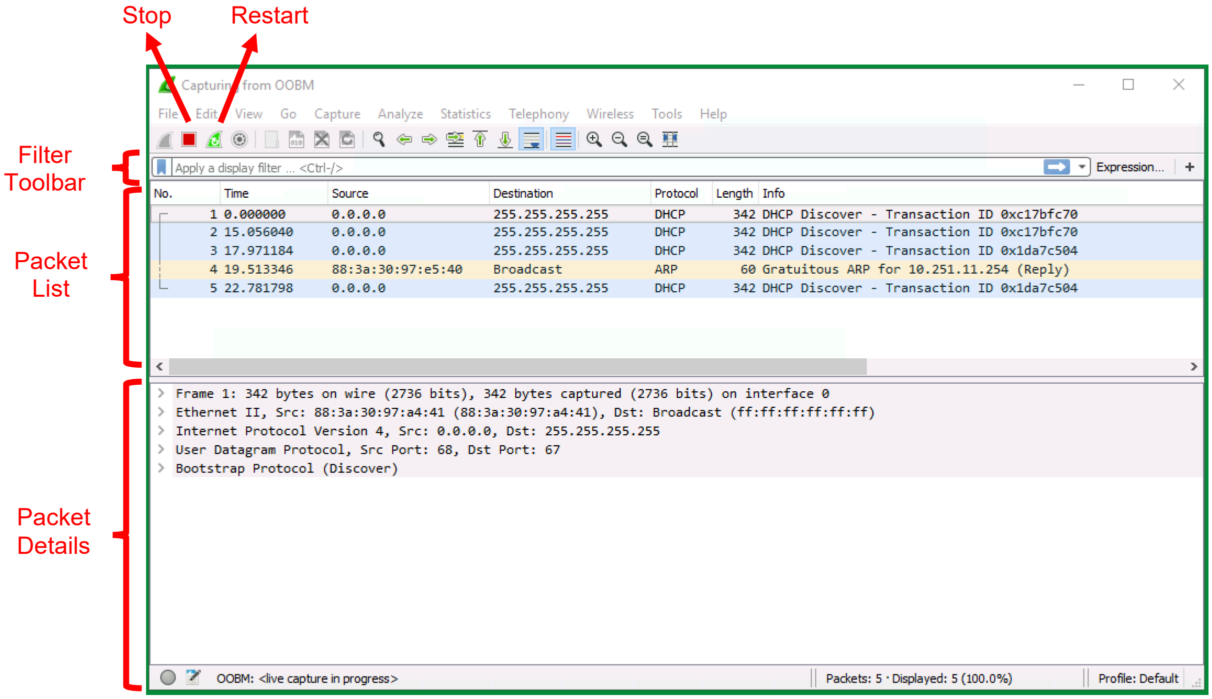
Stop
Restart
Filter Toolbar
Packet List
Packet Details
Capturing from OOBM
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Apply a display filter ... <Ctrl-/>
Expression...

```
No.   Time          Source               Destination        Protocol  Length  Info
    1 0.000000      0.0.0.0              255.255.255.255    DHCP         342  DHCP Discover - Transaction ID 0xc17bfc70
    2 15.056040     0.0.0.0              255.255.255.255    DHCP         342  DHCP Discover - Transaction ID 0xc17bfc70
    3 17.971184     0.0.0.0              255.255.255.255    DHCP         342  DHCP Discover - Transaction ID 0x1da7c504
    4 19.513346     88:3a:30:97:e5:40    Broadcast          ARP           60  Gratuitous ARP for 10.251.11.254 (Reply)
    5 22.781798     0.0.0.0              255.255.255.255    DHCP         342  DHCP Discover - Transaction ID 0x1da7c504
```

```
> Frame 1: 342 bytes on wire (2736 bits), 342 bytes captured (2736 bits) on interface 0
> Ethernet II, Src: 88:3a:30:97:a4:41 (88:3a:30:97:a4:41), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
> Internet Protocol Version 4, Src: 0.0.0.0, Dst: 255.255.255.255
> User Datagram Protocol, Src Port: 68, Dst Port: 67
> Bootstrap Protocol (Discover)
```

OOBM: <live capture in progress> Packets: 5 · Displayed: 5 (100.0%) Profile: Default

Figure 2-5: Wireshark Sections

6. On filter toolbar type “**ip.addr == 10.254.1.22**” with no quotes and hit **[Enter]**. That will instruct Wireshark to only display packets to and from that server.

7. Open a browser and type “**10.254.1.22**” IP address in the URL field and hit **[Enter]**. A page will pop up.

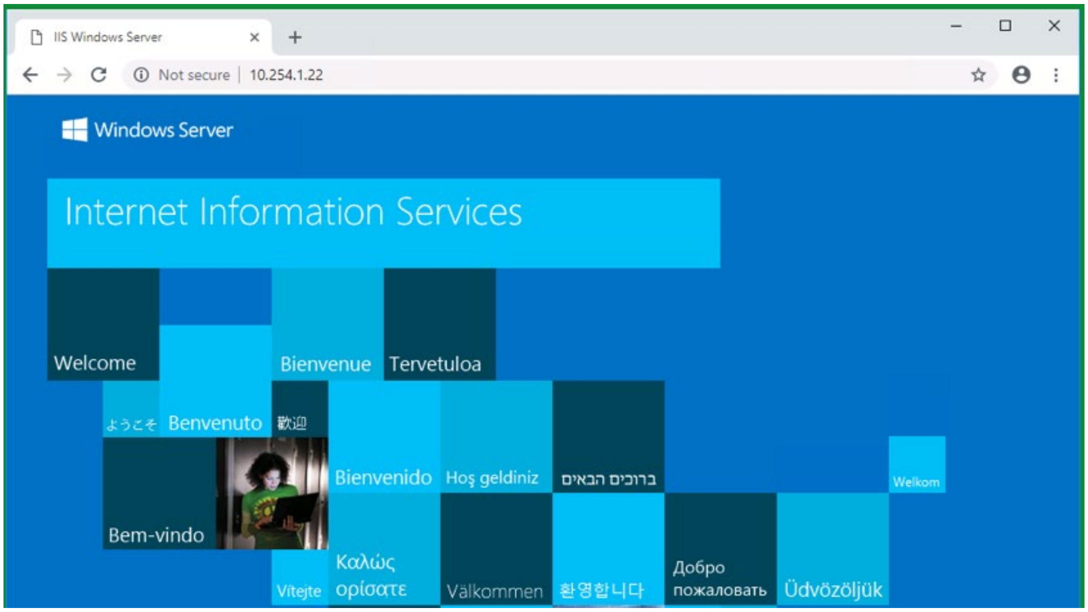

Figure 2-6: Web page

8. Move back to Wireshark. You shall see a long list of entries that represent every single Data Unit exchanged with the server in order to download the page.
9. Scroll all the way up.

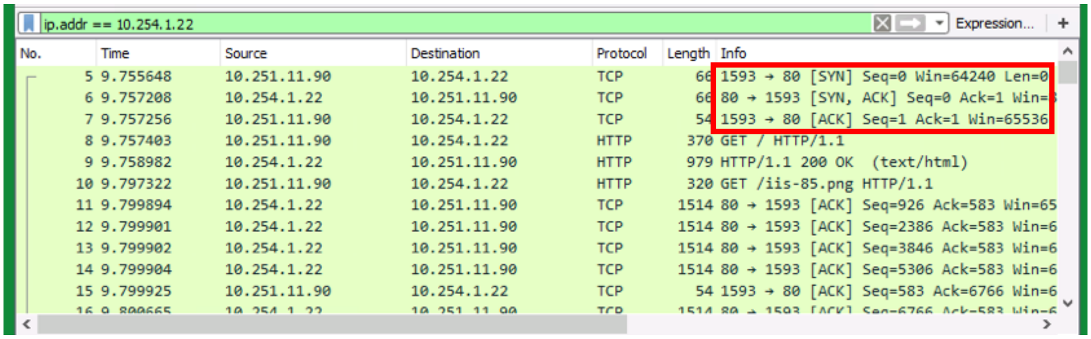

```
ip.addr == 10.254.1.22                                                   Expression...  +
No.   Time        Source          Destination     Protocol  Length Info
   5 9.755648    10.251.11.90    10.254.1.22     TCP         66 1593 → 80 [SYN] Seq=0 Win=64240 Len=0
   6 9.757208    10.254.1.22     10.251.11.90    TCP         66 80 → 1593 [SYN, ACK] Seq=0 Ack=1 Win=
   7 9.757256    10.251.11.90    10.254.1.22     TCP         54 1593 → 80 [ACK] Seq=1 Ack=1 Win=65536
   8 9.757403    10.251.11.90    10.254.1.22     HTTP       370 GET / HTTP/1.1
   9 9.758982    10.254.1.22     10.251.11.90    HTTP       979 HTTP/1.1 200 OK  (text/html)
  10 9.797322    10.251.11.90    10.254.1.22     HTTP       320 GET /iis-85.png HTTP/1.1
  11 9.799894    10.254.1.22     10.251.11.90    TCP       1514 80 → 1593 [ACK] Seq=926 Ack=583 Win=65
  12 9.799901    10.254.1.22     10.251.11.90    TCP       1514 80 → 1593 [ACK] Seq=2386 Ack=583 Win=6
  13 9.799902    10.254.1.22     10.251.11.90    TCP       1514 80 → 1593 [ACK] Seq=3846 Ack=583 Win=6
  14 9.799904    10.254.1.22     10.251.11.90    TCP       1514 80 → 1593 [ACK] Seq=5306 Ack=583 Win=6
  15 9.799925    10.251.11.90    10.254.1.22     TCP         54 1593 → 80 [ACK] Seq=583 Ack=6766 Win=6
  16 9.800665    10.254.1.22     10.251.11.90    TCP       1514 80 → 1593 [ACK] Seq=6766 Ack=583 Win=6
```

Figure 2-7: Data packets

You will first see three packets listed as “SYN”, “SYN, ACK” and “ACK” under the Info column.

What do they mean?

____

____

What are these three packets for?

____

____

10. Select the entry that lists “**GET / HTTP/1.1**” in the Info column. Five entries will appear in the “Packet Details” section including Frame details and Data Link, Network, Transport and Application headers.

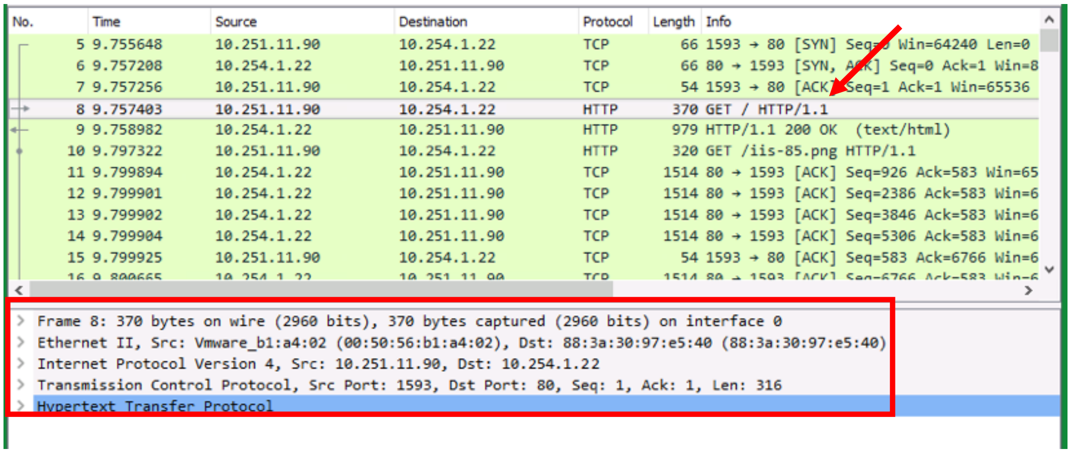

```
No.   Time       Source         Destination    Protocol  Length  Info
    5 9.755648   10.251.11.90   10.254.1.22    TCP         66 1593 → 80 [SYN] Seq=0 Win=64240 Len=0
    6 9.757208   10.254.1.22    10.251.11.90   TCP         66 80 → 1593 [SYN, ACK] Seq=0 Ack=1 Win=8
    7 9.757256   10.251.11.90   10.254.1.22    TCP         54 1593 → 80 [ACK] Seq=1 Ack=1 Win=65536
    8 9.757403   10.251.11.90   10.254.1.22    HTTP       370 GET / HTTP/1.1
    9 9.758982   10.254.1.22    10.251.11.90   HTTP       979 HTTP/1.1 200 OK  (text/html)
   10 9.797322   10.251.11.90   10.254.1.22    HTTP       320 GET /iis-85.png HTTP/1.1
   11 9.799894   10.254.1.22    10.251.11.90   TCP       1514 80 → 1593 [ACK] Seq=926 Ack=583 Win=65
   12 9.799901   10.254.1.22    10.251.11.90   TCP       1514 80 → 1593 [ACK] Seq=2386 Ack=583 Win=6
   13 9.799902   10.254.1.22    10.251.11.90   TCP       1514 80 → 1593 [ACK] Seq=3846 Ack=583 Win=6
   14 9.799904   10.254.1.22    10.251.11.90   TCP       1514 80 → 1593 [ACK] Seq=5306 Ack=583 Win=6
   15 9.799925   10.251.11.90   10.254.1.22    TCP         54 1593 → 80 [ACK] Seq=583 Ack=6766 Win=6
   16 9.800665   10.254.1.22    10.251.11.90   TCP       1514 80 → 1593 [ACK] Seq=6766 Ack=583 Win=6

> Frame 8: 370 bytes on wire (2960 bits), 370 bytes captured (2960 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
> Transmission Control Protocol, Src Port: 1593, Dst Port: 80, Seq: 1, Ack: 1, Len: 316
> Hypertext Transfer Protocol
```

Figure 2-8: Data headers

What protocols are listed in “Frame details” section and what OSI model layers do they belong to?

Data Link header: ____

Network header: ____

Transport header: ____

Application header: ____

11. Click, then expand the “**Ethernet II**” entry.

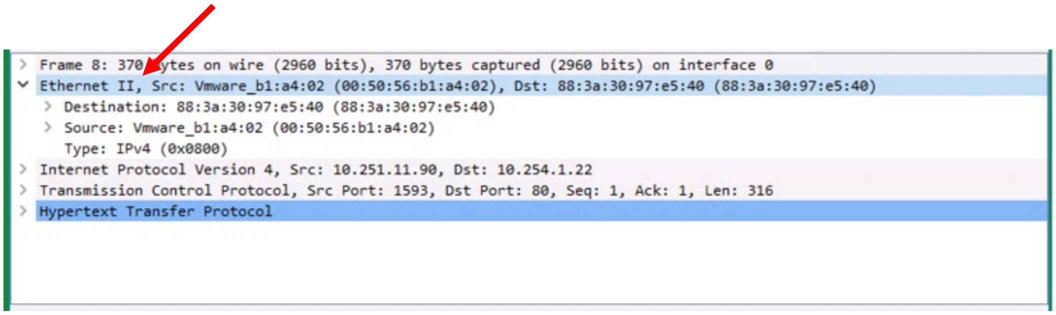

```
> Frame 8: 370 bytes on wire (2960 bits), 370 bytes captured (2960 bits) on interface 0
v Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
  > Destination: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
  > Source: Vmware_b1:a4:02 (00:50:56:b1:a4:02)
    Type: IPv4 (0x0800)
> Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
> Transmission Control Protocol, Src Port: 1593, Dst Port: 80, Seq: 1, Ack: 1, Len: 316
> Hypertext Transfer Protocol
```

Figure 2-9: Data Link layer header

What is the length of the header?

____

What are the values of Destination and Source fields?

____

What is the Type value (also known as Ethertype)?

____

TIP: You can see the header length at the very bottom of the window.

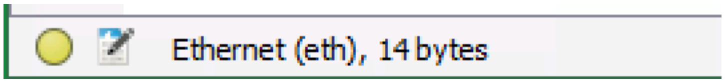

Ethernet (eth), 14 bytes

12. Click, then expand the “**Internet Protocol Version 4**” entry.

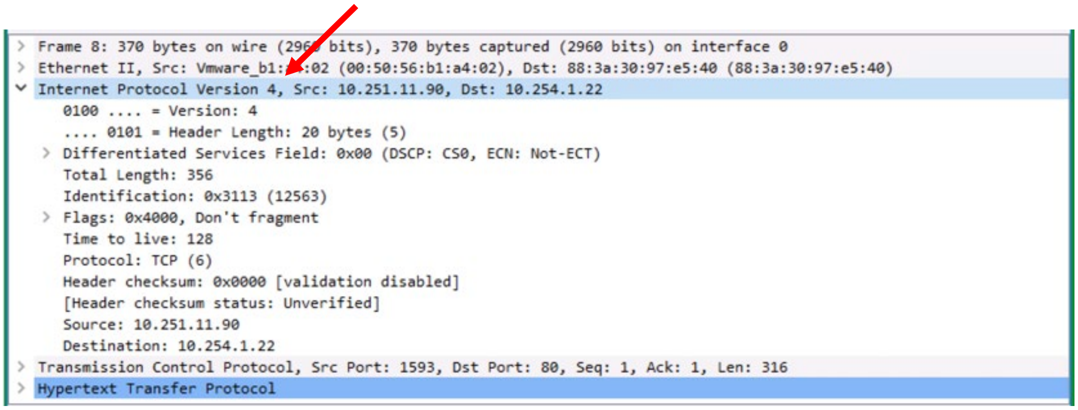

```
> Frame 8: 370 bytes on wire (2960 bits), 370 bytes captured (2960 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
v Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 356
    Identification: 0x3113 (12563)
  > Flags: 0x4000, Don't fragment
    Time to live: 128
    Protocol: TCP (6)
    Header checksum: 0x0000 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.251.11.90
    Destination: 10.254.1.22
> Transmission Control Protocol, Src Port: 1593, Dst Port: 80, Seq: 1, Ack: 1, Len: 316
> Hypertext Transfer Protocol
```

Figure 2-10: Network layer header

What is the length of the header?

What is the protocol version?

What is the Time to live value?

ANSWER: TTL is an 8-bit field with an initial value when the packet is created, every time the packet crosses a layer 3 boundary then TTL is decreased by 1, when it reaches 0 the packet gets discarded.

What is the Protocol number?

What does the IP protocol number represent and what is its main purpose of this field?

ANSWER: IP protocol number or Protocol for short, is a numeric identification of the upper layer protocol contained in the packet's payload. The IANA has assigned unique values to each IP protocol, e.g. ICMP is IP protocol 1, TCP is 6, UDP is 17 and GRE is 47.

What are the values of the Destination and Source fields?

13. Click, then expand the "**Transport Control Protocol**" entry.

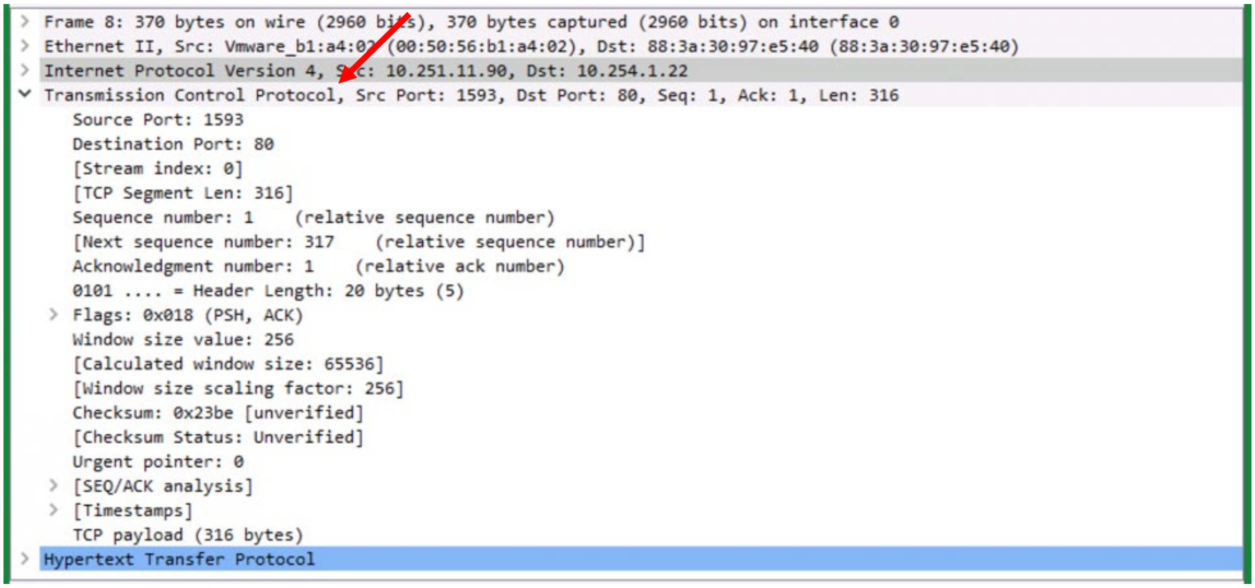

```
> Frame 8: 370 bytes on wire (2960 bits), 370 bytes captured (2960 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
v Transmission Control Protocol, Src Port: 1593, Dst Port: 80, Seq: 1, Ack: 1, Len: 316
     Source Port: 1593
     Destination Port: 80
     [Stream index: 0]
     [TCP Segment Len: 316]
     Sequence number: 1    (relative sequence number)
     [Next sequence number: 317    (relative sequence number)]
     Acknowledgment number: 1    (relative ack number)
     0101 .... = Header Length: 20 bytes (5)
   > Flags: 0x018 (PSH, ACK)
     Window size value: 256
     [Calculated window size: 65536]
     [Window size scaling factor: 256]
     Checksum: 0x23be [unverified]
     [Checksum Status: Unverified]
     Urgent pointer: 0
   > [SEQ/ACK analysis]
   > [Timestamps]
     TCP payload (316 bytes)
> Hypertext Transfer Protocol
```

Figure 2-11: Transport layer header

What is the length of the header?

____

What are the first two fields?

____

____

What are they for?

____

What is the sequence number for?

____

14. Expand “**Flags**”.

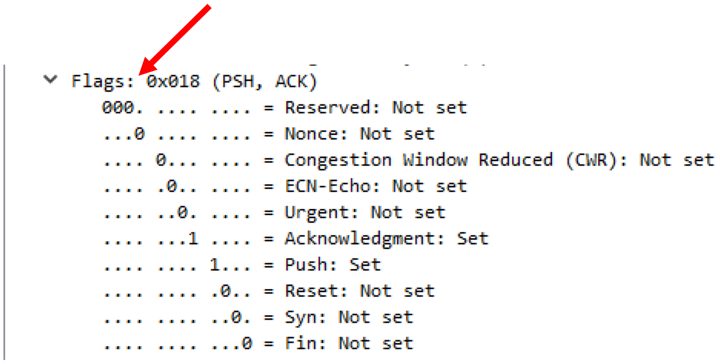

```
v Flags: 0x018 (PSH, ACK)
     000. .... .... = Reserved: Not set
     ...0 .... .... = Nonce: Not set
     .... 0... .... = Congestion Window Reduced (CWR): Not set
     .... .0.. .... = ECN-Echo: Not set
     .... ..0. .... = Urgent: Not set
     .... ...1 .... = Acknowledgment: Set
     .... .... 1... = Push: Set
     .... .... .0.. = Reset: Not set
     .... .... ..0. = Syn: Not set
     .... .... ...0 = Fin: Not set
```

Figure 2-12: TCP flags

Do you know any of them?

Please do some research and find out what the following flags are for?

Acknowledgement:

Reset:

Syn:

Fin:

ANSWER: Flag types are:

- Acknowledgement: Indicates that the acknowledgement field is significant. All packets after the initial SYN packet sent by the client should have this flag set.
- Reset: Reset the connection. Seen on rejected connections.
- Syn: Synchronize the sequence numbers. Seen on new connections.
- Fin: No more data from sender. Seen after a connection is closed.

What is the Window size?

____

What is the Window size for?

____

ANSWER: The window size field is the number of bytes the sender will buffer for the response. During 3-way handshake both sender and receiver will say how large their receive window is.

15. Expand the “**Hypertext Transfer Protocol**” entry.

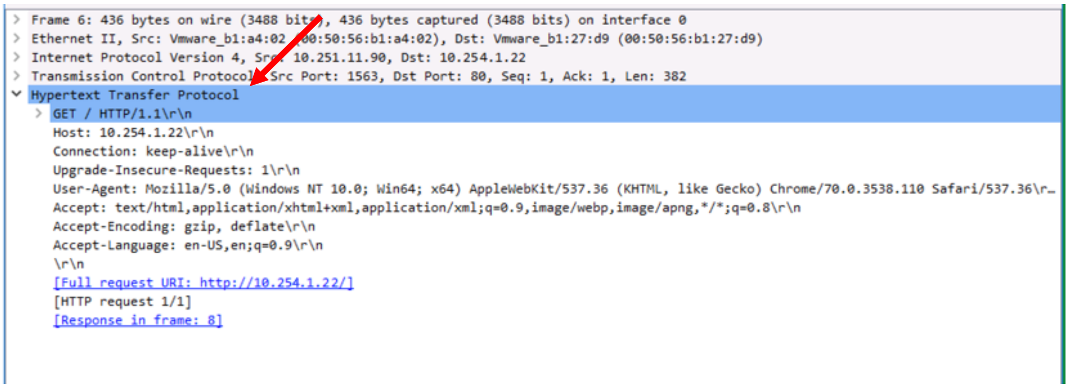

```
> Frame 6: 436 bytes on wire (3488 bits), 436 bytes captured (3488 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: Vmware_b1:27:d9 (00:50:56:b1:27:d9)
> Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
> Transmission Control Protocol, Src Port: 1563, Dst Port: 80, Seq: 1, Ack: 1, Len: 382
v Hypertext Transfer Protocol
  > GET / HTTP/1.1\r\n
    Host: 10.254.1.22\r\n
    Connection: keep-alive\r\n
    Upgrade-Insecure-Requests: 1\r\n
    User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/70.0.3538.110 Safari/537.36\r…
    Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,image/apng,*/*;q=0.8\r\n
    Accept-Encoding: gzip, deflate\r\n
    Accept-Language: en-US,en;q=0.9\r\n
    \r\n
    [Full request URI: http://10.254.1.22/]
    [HTTP request 1/1]
    [Response in frame: 8]
```

Figure 2-13: Application Header

IMPORTANT: In Hyper Transfer Protocol or HTTP’s header there are 4 main commands: GET, PUSH, PUT and DELETE. Usually after the 3-way handshake, the first HTTP payload has a GET instruction in order to download the web page.

After requesting the web page, there will be a lot of packets coming from the server. These are acknowledged by the client and displayed as the black with red entries (image below), they contain the web page itself. Once the page is fully loaded in the browser there is a FIN segment coming from the client signaling the end of the session. It is followed by similar one from the server, and finally a last ACK is sent by the client.

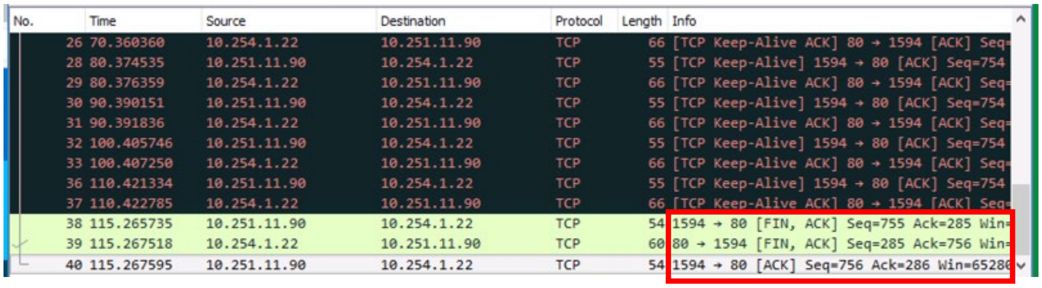

No.	Time	Source	Destination	Protocol	Length	Info
26	70.360360	10.254.1.22	10.251.11.90	TCP	66	[TCP Keep-Alive ACK] 80 → 1594 [ACK] Seq=
28	80.374535	10.251.11.90	10.254.1.22	TCP	55	[TCP Keep-Alive] 1594 → 80 [ACK] Seq=754
29	80.376359	10.254.1.22	10.251.11.90	TCP	66	[TCP Keep-Alive ACK] 80 → 1594 [ACK] Seq=
30	90.390151	10.251.11.90	10.254.1.22	TCP	55	[TCP Keep-Alive] 1594 → 80 [ACK] Seq=754
31	90.391836	10.254.1.22	10.251.11.90	TCP	66	[TCP Keep-Alive ACK] 80 → 1594 [ACK] Seq=
32	100.405746	10.251.11.90	10.254.1.22	TCP	55	[TCP Keep-Alive] 1594 → 80 [ACK] Seq=754
33	100.407250	10.254.1.22	10.251.11.90	TCP	66	[TCP Keep-Alive ACK] 80 → 1594 [ACK] Seq=
36	110.421334	10.251.11.90	10.254.1.22	TCP	55	[TCP Keep-Alive] 1594 → 80 [ACK] Seq=754
37	110.422785	10.254.1.22	10.251.11.90	TCP	66	[TCP Keep-Alive ACK] 80 → 1594 [ACK] Seq=
38	115.265735	10.251.11.90	10.254.1.22	TCP	54	1594 → 80 [FIN, ACK] Seq=755 Ack=285 Win=
39	115.267518	10.254.1.22	10.251.11.90	TCP	60	80 → 1594 [FIN, ACK] Seq=285 Ack=756 Win=
40	115.267595	10.251.11.90	10.254.1.22	TCP	54	1594 → 80 [ACK] Seq=756 Ack=286 Win=6528

Task 2: UDP header

Objectives

Now you will look into a UDP header and compare it with the TCP one.

Steps

PC-1

1. Click the restart button then click “**Continue without Saving**” button. This will clean up the packet capture.

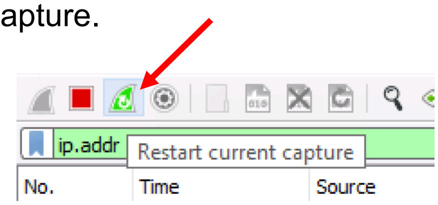

Figure 2-14: Wireshark restart

2. Open **3CDaemon**, there should be a shortcut on the Desktop.

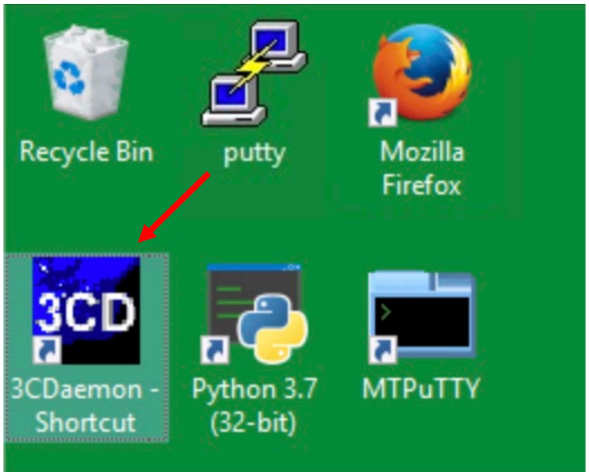

Figure 2-15: 3CDaemon

3. Click on the “**Tftp Client**” tab.
4. For TFTP Server Address type “**10.254.1.22**”
5. On Operation select “**Receive File**”.
6. For Remote File Name type CXF.txt.
7. Click the “**…**” button next to “**Local File Name**” field, then select Desktop as destination directory and type **CXF.txt** as file name.
8. Click **Save** button.

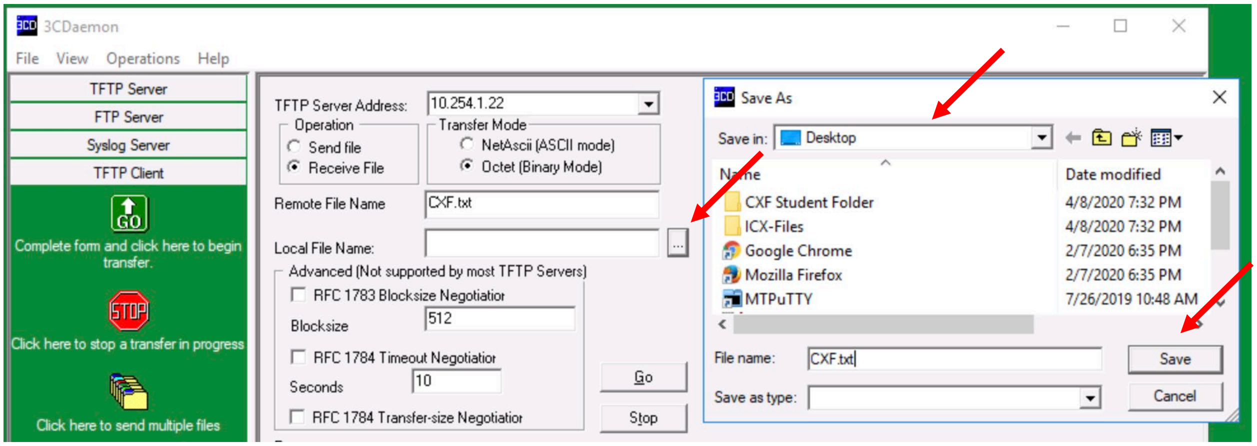
3CDaemon
File View Operations Help
TFTP Server
FTP Server
Syslog Server
TFTP Client
GO
Complete form and click here to begin transfer.
STOP
Click here to stop a transfer in progress
Click here to send multiple files
Go
Stop
Save As
Name
Date modified
Save
Cancel

TFTP Server Address: 10.254.1.22

Operation: ○ Send file ● Receive File

Transfer Mode: ○ NetAscii (ASCII mode) ● Octet (Binary Mode)

Remote File Name: CXF.txt

Local File Name:

Advanced (Not supported by most TFTP Servers)

□ RFC 1783 Blocksize Negotiation

Blocksize: 512

□ RFC 1784 Timeout Negotiation

Seconds: 10

□ RFC 1784 Transfer-size Negotiation

Save in: Desktop

Name	Date modified
CXF Student Folder	4/8/2020 7:32 PM
ICX-Files	4/8/2020 7:32 PM
Google Chrome	2/7/2020 6:35 PM
Mozilla Firefox	2/7/2020 6:35 PM
MTPuTTY	7/26/2019 10:48 AM

File name: CXF.txt

Save as type:

Figure 2-16: 3CDaemon – Local File

9. Back in TFTP Client click the **Go** button. The software will begin a TFTP connection and download the file.

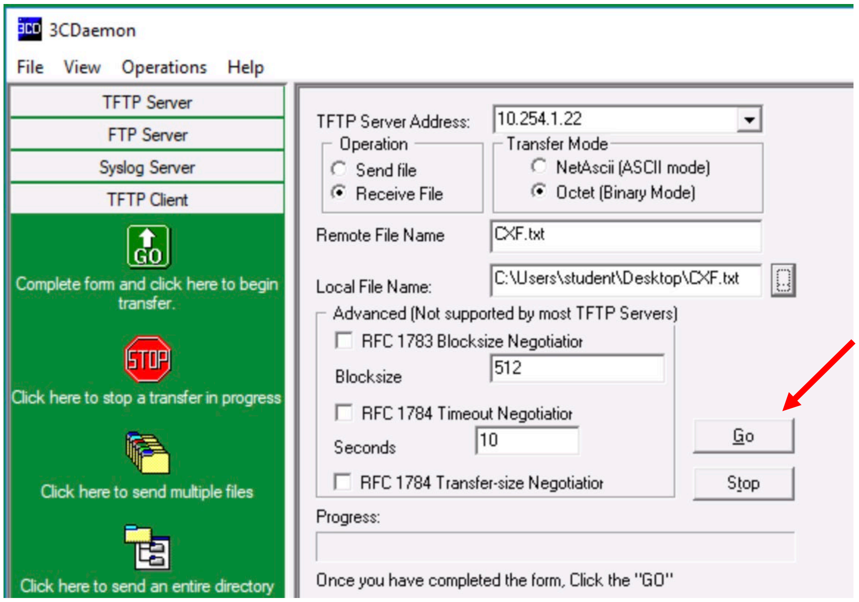
3CDaemon
File View Operations Help
TFTP Server
FTP Server
Syslog Server
TFTP Client
GO
Complete form and click here to begin transfer.
STOP
Click here to stop a transfer in progress
Click here to send multiple files
Click here to send an entire directory
Go
Stop

TFTP Server Address: 10.254.1.22

Operation: ○ Send file ● Receive File

Transfer Mode: ○ NetAscii (ASCII mode) ● Octet (Binary Mode)

Remote File Name: CXF.txt

Local File Name: C:\Users\student\Desktop\CXF.txt

Advanced (Not supported by most TFTP Servers)

□ RFC 1783 Blocksize Negotiation

Blocksize: 512

□ RFC 1784 Timeout Negotiation

Seconds: 10

□ RFC 1784 Transfer-size Negotiation

Progress:

Once you have completed the form, Click the "GO"

Figure 2-17: TFTP client settings

10. Move to **Wireshark**. You will see a new capture with all packets involved in the transfer.

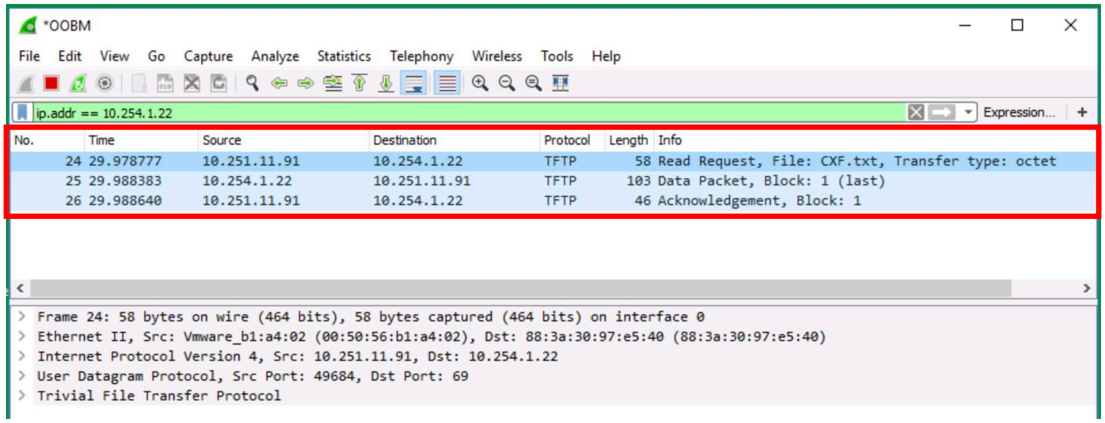
*OOBM
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
ip.addr == 10.254.1.22
Expression...

```
No.  Time        Source        Destination   Protocol Length Info
  24 29.978777   10.251.11.91  10.254.1.22   TFTP       58 Read Request, File: CXF.txt, Transfer type: octet
  25 29.988383   10.254.1.22   10.251.11.91  TFTP      103 Data Packet, Block: 1 (last)
  26 29.988640   10.251.11.91  10.254.1.22   TFTP       46 Acknowledgement, Block: 1
```

```
> Frame 24: 58 bytes on wire (464 bits), 58 bytes captured (464 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.91, Dst: 10.254.1.22
> User Datagram Protocol, Src Port: 49684, Dst Port: 69
> Trivial File Transfer Protocol
```

Figure 2-18: TFTP traffic capture

Is there any Three-way handshake session establishment?

____

11. Click the first packet (Read Request).
12. Select and expand the "**User Datagram Protocol**" entry in the Packet Details section.

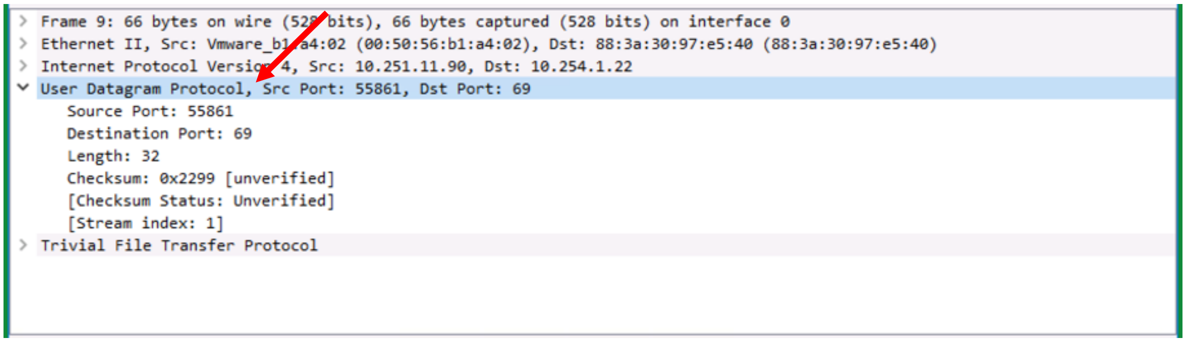

```
> Frame 9: 66 bytes on wire (528 bits), 66 bytes captured (528 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.90, Dst: 10.254.1.22
v User Datagram Protocol, Src Port: 55861, Dst Port: 69
     Source Port: 55861
     Destination Port: 69
     Length: 32
     Checksum: 0x2299 [unverified]
     [Checksum Status: Unverified]
     [Stream index: 1]
> Trivial File Transfer Protocol
```

Figure 2-19: User Datagram Protocol

What is the length of the header?

____

What is the first impression when comparing with the TCP header (Task 1 step 13)?

____

____

What fields do they have in common?

____

Can you see any Acknowledgment flag embedded in the header?

____

13. Click and expand the “**Trivial File Transfer Protocol**” entry.

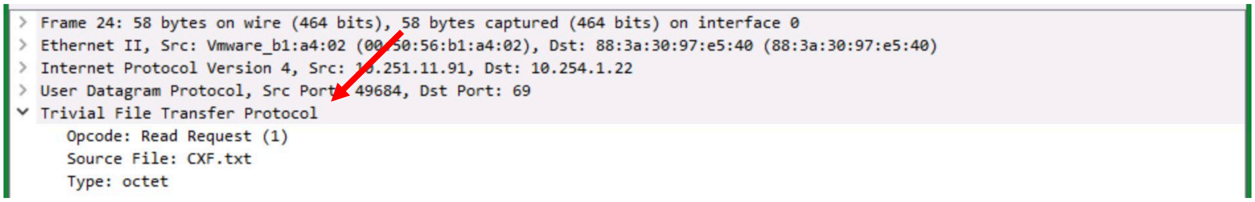

```
> Frame 24: 58 bytes on wire (464 bits), 58 bytes captured (464 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.91, Dst: 10.254.1.22
> User Datagram Protocol, Src Port: 49684, Dst Port: 69
v Trivial File Transfer Protocol
    Opcode: Read Request (1)
    Source File: CXF.txt
    Type: octet
```

Figure 2-20: TFTP traffic capture

NOTE: This is the TFTP application header, just by looking in its contents you can tell this is the CXF.txt file request sent by the client.

14. Click the last packet (**Acknowledgement**). It will automatically show the TFTP header contents

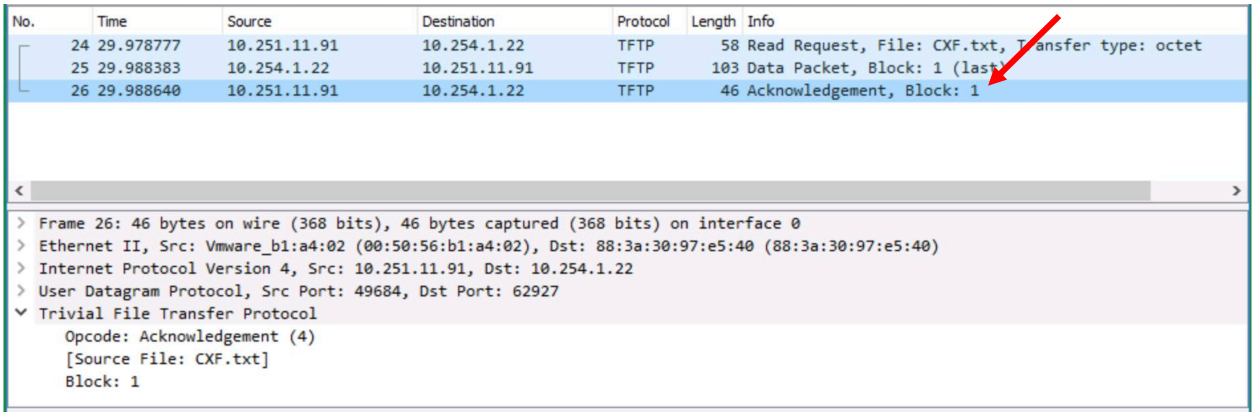

```
No.   Time        Source         Destination    Protocol  Length  Info
   24 29.978777   10.251.11.91   10.254.1.22    TFTP          58  Read Request, File: CXF.txt, Transfer type: octet
   25 29.988383   10.254.1.22    10.251.11.91   TFTP         103  Data Packet, Block: 1 (last)
   26 29.988640   10.251.11.91   10.254.1.22    TFTP          46  Acknowledgement, Block: 1

> Frame 26: 46 bytes on wire (368 bits), 46 bytes captured (368 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a4:02 (00:50:56:b1:a4:02), Dst: 88:3a:30:97:e5:40 (88:3a:30:97:e5:40)
> Internet Protocol Version 4, Src: 10.251.11.91, Dst: 10.254.1.22
> User Datagram Protocol, Src Port: 49684, Dst Port: 62927
v Trivial File Transfer Protocol
    Opcode: Acknowledgement (4)
    [Source File: CXF.txt]
    Block: 1
```

Figure 2-21: TFTP Traffic Capture

What is the Opcode field value?

____

IMPORTANT: Due the lack of acknowledgement at the transport level, some UDP based applications do support the feature at Layer 7 level, this is the case of TFTP.

Also notice how, unlike TCP, the transmission suddenly stops without any FIN signaling at the transport layer. This is because at the application layer level the TFTP server told the client how many bytes the file has, once those bytes were sent and acknowledged (again at Layer 7), then both parties assume the session is over.

You have completed Lab 2!

AOS-CX Switching Fundamentals

Lab 3: Initial Setup

Overview

BigStartup is a small business that just started operations a few months ago. The owners have determined the need to rent a small portion of a nearby building's floor (The East Wing) from Cheap4Rent Properties in order to house a new group of employees they just hired. These employees will be using Windows PCs and will have a few networking connectivity requirements in their daily operations, such as printing and file sharing. Because of this, you have been contacted to provide network consulting services, as well as take care of configuring and managing the switching equipment that BigStartup recently purchased.

Objectives

After completing this lab, you will be able to:

- Set your gear in factory values
- Navigate through the AOS-CX command line interface (CLI)
- Define a hostname on 6300-A switch
- Disable unused interfaces
- Save device's configuration and create checkpoints

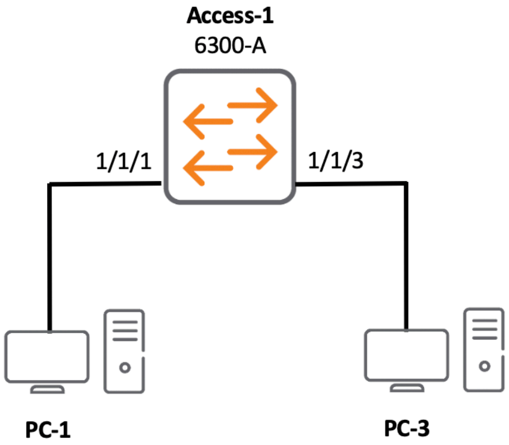

Figure 2-1: Lab Topology

Task 1: Explore the AOS-CX Switch CLI

Objectives

In this task, you will explore and become more familiar with the AOS-CX switch CLI. Do not be afraid to try out different commands on the CLI: you will learn by experimenting!

Steps

6300-A

1. Open a console connection to the **6300-A**. Login using **admin** and **no password**.
2. Hit the **[?]** key to show the available commands that you can execute in the current command context.

```
6300# ?
  auto-confirm  Disables user confirmation, and executes the operation without
                prompting
  boot          Reboot all or part of the system; configure default boot
                parameters
  checkpoint    Checkpoint information
  clear         Reset functions
  configure     Configuration from vty interface
  copy          Copy data or files to/from the switch
  debug         Configure debug logging
  diagnostics   Change diagnostic commands availability
  disable       Turn off privileged mode command
  end           End current mode and change to enable mode
  erase         Erase device information or files
  exit          Exit current mode and change to previous mode
  https-server  HTTPS Server management
  list          Print command list
  member        VSF member selection
  no            Negate a command or set its defaults
  page          Enable page break
  ping          Send IPv4 ping requests to a device on the network
  ping6         Send IPv6 ping requests to a device on the network
  port-access   Port based network access.
  repeat        Repeat a list of commands from history
  show          Show running system information
  ssh           Configure SSH.
  start-shell   Start Bash shell
  top           Top command
  traceroute    Trace the IPv4 route to a device on the network
  traceroute6   Trace the IPv6 route to a device on the network
  usb           Commands to control the USB Port
  vsf           Virtual Switching Framework (VSF) commands
  write         Write running configuration to memory, network, or terminal
```

TIP: Page through the commands available at this level. Some important commands available at this level include.

- **show**, which enables you to examine current configuration parameters
- **copy**, which enables you to back up the switch configuration
- **ping** and **traceroute**, which are connectivity test tools

3. List the parameters available for the **show** command. By typing “**show”** followed by **[?].**

```
6300# show ?
  aaa                     Authentication, Authorization and Accounting.
  access-list             Access control list (ACL)
  accounting              Show local accounting information
  active-gateway          Show active gateway settings
  alias                   Short names configured for a set of commands
  allow-unsafe-updates    Show allowed non-failsafe updates
  arp                     Show IPv4 addresses from neighbor table
  aruba-central           Configure Aruba-Central
  banner                  Show one of the configured system banners
  bfd                     BFD information
  bgp                     BGP specific commands
  bluetooth               Display information about Bluetooth wireless management
  boot-history            Display boot history details
  capacities              Show system capacities and its values.
  capacities-status       Show system capacities status and its values.
  cdp                     Show various CDP settings
  checkpoint              Checkpoint information
                             ←---- output omitted ---→
```

4. Scroll through.
5. Type “**disable”**.

```
6300# disable
6300>
```

How has the prompt changed?

____

ANSWER: This turns privileged mode off, which means only basic commands with no control upon the device will be available.

6. Hit the **[?]** key to show the available commands that you can execute in this non-Privileged command context.

```
6300> ?
  clear   Reset functions
  enable  Turn on privileged mode command
  exit    Exit current mode and change to previous mode
  list    Print command list
  no      Negate a command or set its defaults
  page    Enable page break
  repeat  Repeat a list of commands from history
  show    Show running system information
  top     Top command
  user    User account
                            ←---- output omitted ---→
```

IMPORTANT: Available commands in both privileged and no privileged modes are different, this is used as a basic role-based access control for defining what operators can do when logged into the device.

7. Type **"enable"** and hit enter, this will turn privileged mode back again.

```
6300> enable
6300#
```

8. Type "**co**" then hit the [**tab**] key twice to list commands that start with "co":

```
6300# co[tab][tab]
```

What does the CLI display?

9. Type “**con**” followed by a single **[tab]** hit.

```
6300# configure
```

What has just happened to the command?

TIP: You can execute any command as soon as you have entered an unambiguous character string. For instance, conf [Enter] will have the same effect as configure [Enter].

10. Hit **[Enter]** key. This takes you to global configuration mode, where you can start making changes that take immediate effect upon the device’s configuration.

```
6300# configure
6300(config)#
```

11. Hit **[?]** key to show the available commands that you can execute in the global config mode.

```
6300(config)# ?
  aaa                    Configure Authentication, Authorization and Accounting
                         feature.
  access-list            Access control list (ACL)
  alias                  Create a short name for the specified command(s)
  allow-unsafe-updates   Allow non-failsafe updates of programmable devices
  apply                  Apply a configuration record
  aruba-central          Configure Aruba-Central
  banner                 Customize login banner
  bfd                    Enable Bidirectional Forwarding Detection (BFD)
  bluetooth              Configure Bluetooth wireless management
  cdp                    Configure CDP operating mode
  checkpoint             Configure checkpoint related feature
  class                  Configure classifier class
  cli-session            Configure CLI session management
                             ←---- output omitted ---→
```

NOTE: You can notice how commands available here are different than in previous CLI modes due the configuration nature of them.

12. Type **“interface 1/1/1”** then hit **[enter]**. You will be moved to the interface sub configuration mode.

```
6300(config)# interface 1/1/1
6300(config-if)#
```

13. Hit **[?]** key. Again, you will see a different list of available commands for this sub context.

```
6300(config-if)# ?
  aaa                Configure Authentication, Authorization and Accounting
feature.
  apply              Apply a configuration record
  arp                Configure ARP commands
  bfd                Set BFD configuration
  cdp                Configure CDP operating mode
  description        Add an interface description
  dhcpv4-snooping    Configure DHCPv4-Snooping
  dhcpv6-snooping    Configure DHCPv6-Snooping
  end                End current mode and change to enable mode
  exit               Exit current mode and change to previous mode
  flow-control       Configure flow control
```

14. Type “**end**” and hit **[Enter]**.

```
6300(config-if)# end
6300#
```

What has just happened to the command prompt?

____

____

Next, you will enter a command that is invalid, and then fix issues with it by using the command-recall feature.

15. Enter this command exactly as shown: "**show hitory**".

```
6300# show hitory
Invalid input: hitory
```

16. Recall the command by pressing the **[Up]** arrow key.
17. Go to the beginning of the command with the **[CTRL][a]** shortcut.
18. Go to the end of the command line with the **[CTRL][e]** shortcut.
19. With the **[Left]** and **[Right]** arrow keys, move your cursor to the correct position in "**hitory**" and place the letter "s".
20. Press the **[Enter]** key at any time (no matter where your cursor is) to execute the command.

TIP: Repeating commands can be a useful way to enter similar commands more quickly, as well as to correct mistakes in commands.

```
6300# show history
6     disable
5     enable
4     configure
3     interface 1/1/1
2     end
1     show hitory
6300#
```

21. Recall the wrong command by pressing the **[Up]** arrow key **two times**.

IMPORTANT: Using the [CTRL][w] shortcut for removing the word that is preceding the cursor is useful in cases you want to either quickly correct a typo or you intend to use another form of the root command.

```
6300# show hitory [CTRL][w]
6300# show
```

22. Add “**system**” to the show command followed by “**?**”.

```
6300# show system ?
  resource-utilization  Utilization metrics of various system resources
  serviceos             Display serviceOS information
  <cr>
6300# show system
```

What options are available for the “show system” command?

____

NOTE: Notice the **<cr>** at the end, this means that you can execute the command without supplying any further parameters.

23. View the system resource utilization on the switch.

```
6300# show system resource-utilization


                            ←---- output omitted ---→
dhcp-server-ada             0                1               8
systemd-udevd               0                0              14
hpe-restd                   0                1              14
hpe-entityd                 0                0              10
ata_sff                     0                0               0
mmcqd/0                     0                0               0
bled                        0                0              12
mmcqd/0rpmb                 0                0               0
vlanremapd                  0                0               9
powerd                      0                0              12
rpciod                      0                0               0
nfsd                        0                0               0
rcu_sched                   0                0               0
```

TIP: You will notice that a long output automatically populates overrunning on the screen, not giving you the chance to read the first lines. You can use the “page” command for displaying subsequent command outputs in portions and giving you the ability to control when to display the next page by hitting the space bar.

24. Use the “**page**” command followed by “**system resource-utilization**”.

```
6300# page
6300# show system resource-utilization

System Resources:
Processes: 202
CPU usage(%): 1
Memory usage(%): 16
Open FD's: 6048
mmc-type-a: Endurance = 0-10%, Health = normal
mmc-type-b: Endurance = 0-10%, Health = normal

ProcessCPU Usage(%)Memory Usage(%)  Open FD's
--------------------------------------------------------------------------
kworker/u8:2               0                0                0
tacacs-srv-trkd            0                0                9
acctsyslogd                0                0                6
watchdogd                  0                0                0
prometheus                 0                1               11
nginx                      0                0               13
kblockd                    0                0                0
ovsdb-server              14                0              125
migration/3                0                0                0
fpsLink1                   0                0                0
kauditd                    0                0                0
mmcqd/0gp1                 0                0                0
 -- MORE --, next page: Space, next line: Enter, quit: q
```

What has changed in this new output?

____

____

ANSWER: The command shows the current CPU and memory utilization of the system and the per process utilization.

What is current CPU and Memory utilization of the switch?

TIP: Alternately you can use the "top cpu" and "top memory" commands for displaying these numbers. A key difference between "show system resource-utilization" and "top" commands is that "top" commands list higher resource using commands first. Also, the output displays the processes' ID, status and the user that is running the command (the system or a real user logged into the device).

NOTICE: High CPU utilization is a symptom of an unstable process or situation happening in the system, such a layer 2, layer 3 or layer 7 loop.

25. Hit **[space]** a few times to scroll all the way down or **[q]** key.

26. Try "**show system**" command. This version of the command will also show current hostname, description SNMP contact and location, serial number, base MAC address, up time, etc.

```
6300# show system
Hostname             : 6300
System Description : FL.10.04.0030
System Contact       :
System Location      :

Vendor               : Aruba
Product Name         : JL668A 6300F 24G 4SFP56 Sw
Chassis Serial Nbr : SG90KN70HX
Base MAC Address     : 883a30-983000
AOS-CX Version : FL.10.04.0030

Time Zone            : UTC

Up Time              : 2 hours under a minute
CPU Util (%)         : 5
Memory Usage (%)     : 16
```

What is current Hostname?

What is Chassis serial number?

____

What is system base MAC address?

____

What is system Up Time?

____

27. Execute the “**list**” command.

```
6300# list
  show hostname
  show domain-name
  list
  configure { terminal }
  disable
  exit
  end
  page
  page <2-1000>
  no page
  show running-config
  show session-timeout
  start-shell
  auto-confirm
  no auto-confirm
  diagnostics
  no diagnostics
  show history {timestamp}
  repeat { id <A:1-500>|count <1-1000>|delay <1-1000> }
  show vrf
  show vrf VRF
  show dhcp client vendor-class-identifier
  show ztp information
 -- MORE --, next page: Space, next line: Enter, quit: q
```

What does the output display?

____

____

____

IMPORTANT: “list” command shows the right syntax for all commands available at the current context along with their variants and extensions. This can be helpful for discovering new commands and previewing their different forms.

28. Execute the “**show version**” command.

```
6300# show version
-----------------------------------------------------------------------------
AOS-CX
(c) Copyright 2017-2019 Hewlett Packard Enterprise Development LP
-----------------------------------------------------------------------------
Version      : FL.10.04.0030
Build Date   : 2019-11-15 10:37:55 PST
Build ID     : AOS-CX:FL.10.04.0030:ff84f1ebd5b2:201911151752
Build SHA    : ff84f1ebd5b2765e65cc6de982b8ce8d16228050
Active Image : primary

Service OS Version : FL.01.05.0003
BIOS Version       : FL.01.0002
6300#
```

What main AOS-CX code version is running in the system?

____

29. Execute the “**show images**” command.

```
6300# show images
---------------------------------------------------------------------------
AOS-CX Primary Image
---------------------------------------------------------------------------
Version : FL.10.04.0030
Size    : 722 MB
Date    : 2019-11-15 10:37:55 PST
```

```
SHA-256 : abe5ec454be522ce6e3947db1c09fbdb2bfe72ae0447f5055ad592dadc422deb

---------------------------------------------------------------------------
AOS-CX Secondary Image
---------------------------------------------------------------------------
Version : FL.10.04.0030
Size    : 722 MB
Date    : 2019-11-15 10:37:55 PST
SHA-256 : abe5ec454be522ce6e3947db1c09fbdb2bfe72ae0447f5055ad592dadc422deb

Default Image : primary

------------------------------------------------------
Management Module 1/1 (Active)
------------------------------------------------------
Active Image       : primary
Service OS Version : FL.01.05.0003
BIOS Version       : FL.01.0002

6300#
```

How many images does the system support?

What is the default image?

30. Execute the "**show capacities**" command (be prepared for a long output).

```
6300# show capacities

System Capacities:
Capacities Name                                                          Value
---------------------------------------------------------------------------------
                            ←---- output omitted ---→
Maximum number of entries in an Access Control List                       8000
Maximum number of entries in a class                                      1000
Maximum number of entries in an Object Group                                64
Maximum number of entries in a policy                                      128
                            ←---- output omitted ---→
Maximum number of IP neighbors (IPv4+IPv6) that can be configured on the
system                                                                   49152
Maximum number of IP source lockdown bindings allowed on the system      16384
Maximum number of GRE IPv4, "IPv6 in IPv4" and "IPv6 in IPv6" tunnels in
a system                                                                   127
```

```
Maximum number of IPv4 neighbors to hold in ARP cache before performing garbage
collection                                                                 262144
Maximum number of IPv4 neighbors that can be configured on the system       49152
Maximum number of IPv6 neighbors in hold in ARP cache before performing
garbage collection                                                         131072
Maximum number of IPv6 neighbors that can be configured on the system       49152
Maximum number of L2 MAC addresses supported in the system                  32768
Maximum number of L3 Groups for IP Tunnels and ECMP Groups                   2047
Maximum number of L3 Destinations for Routes, Nexthops in Tunnels and
ECMP groups                                                                  2045
Maximum number of configurable LAG ports                                      256
Maximum number of members supported by a LAG port                               8
                           ←---- output omitted ---→
Maximum number of unique tunnel TTLs in a system                                4
Maximum number of routes (IPv4+IPv6) on the system                          65536
Maximum number of IPv4 routes on the system                                 65536
Maximum number of IPv6 routes on the system                                 61440
Maximum number of VLANs supported in the system                              4094
Maximum number of VLAN Translation rules supported                           4000
                           ←---- output omitted ---→
```

What is the maximum amount of access control entries per Access-list supported in the system?

What is the maximum amount of MAC addresses supported in the system?

What is the maximum amount of IP routes (IPv4 and IPv6 combined) supported in the system?

What is the maximum amount of VLANs supported in the system?

TIP: A similar command: "**show capacities-status**" displays similar information plus the amount of resources/entries already consumed by the current device state.

31. Execute the "**show interface 1/1/1**" command.

IMPORTANT: Output displays among many things, the interface state, interface type, current speed and duplex settings, MTU configured, port VLAN mode: access or trunk, and interface counters.

```
6300# show interface 1/1/1

Interface 1/1/1 is up
 Admin state is up
 Link transitions: 1
 Description:
 Hardware: Ethernet, MAC Address: 88:3a:30:98:30:27
 MTU 1500
 Type 1GbT
 Full-duplex
 qos trust none
 Speed 1000 Mb/s
 Auto-negotiation is on
 Flow-control: off
 Error-control: off
 MDI mode: MDIX
 VLAN Mode: access
 Access VLAN: 1
 Rx
            22 input packets            4695 bytes
             0 input error                22 dropped
             0 CRC/FCS
 Tx
           352 output packets          46087 bytes
             7 input error                 7 dropped
             0 collision
6300#
```

What is the interface type?

32. Now try "**show interface 1/1/28**" command.

```
6300# show interface 1/1/28

Interface 1/1/28 is up
 Admin state is up
 Link transitions: 3
 Description:
 Hardware: Ethernet, MAC Address: 88:3a:30:98:30:0d
 MTU 1500
 Type SFP+DAC1
 Full-duplex
 qos trust none
 Speed 10000 Mb/s
 Auto-negotiation is off
 Flow-control: off
 Error-control: off
 VLAN Mode: access
 Access VLAN: 1
 Rx
         37775 input packets         5124978 bytes
             0 input error             40317 dropped
             0 CRC/FCS
 Tx
          2216 output packets         279307 bytes
            10 input error                10 dropped
             0 collision
6300#
```

What is the interface type?

____

ANSWER: Interfaces 1/1/25 to 1/1/28 in a 24 ports switch model and 1/1/49 to 1/1/52 in a 48 ports switch model are SPF+ 25Gig capable interfaces that support either transceivers or Direct Attached Cables (DACs). In this case port 28 has a 10Gig DAC attached.

33. Execute the "**show interface transceiver**" command.

```
6300# show interface transceiver
---------------------------------------------------------------
Port      Type           Product   Serial          Part
                         Number    Number          Number
---------------------------------------------------------------
1/1/25    SFP+DAC1       J9281D    CN97KBZ55Y      8121-1300
```

```
1/1/26    SFP+DAC1       J9281D    CN97KBZ46K      8121-1300
1/1/27    SFP+DAC1       J9281D    CN94KBZ97T      8121-1300
1/1/28    SFP+DAC1       J9281D    CN94KBZ9RZ      8121-1300
```

Task 2: Configure Initial Settings

Objectives

In this task, you will explore the AOS-CX configuration script and make minor customization changes like setting a hostname, setting interface descriptions and disabling unused ports. Also, you will ask the system to display the event log contents.

Steps

6300-A

1. Open a console connection to the 6300-A. Login using **admin** and **no password**.
2. Issue the "**show running-config**" command to display the current configuration of the system.

> **NOTE:** You will notice that most portions of the configuration are shown by listing the switch ports and their settings. The code version and actual admin account are listed first.

```
6300# show running-config
Current configuration:
!
!Version AOS-CX FL.10.04.0030
!export-password: default
user admin group administrators password ciphertext
AQBapRp8LdLe6JVQcygkkgkv/oylHnTJxLb0P2kWEByJaxIhYgAAAMYWlwysPrGTTG/dRu2WG5zqNV4Oi
Nx7ZRZiyhYyJ1T0T4dW3yFaOcCbxCb4qVq1gdlHi
4qdcjN2ILQQyQi39b5rVsZSv1LjaoluZnvyvyZTC2+kN7fTKxxl16Nf8Fq6T+8I
!
!
ssh server vrf default
ssh server vrf mgmt
!
!
vlan 1
no spanning-tree
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    no routing
    vlan access 1
interface 1/1/2
    no shutdown
    no routing
    vlan access 1
                           ←---- output omitted ---→
interface 1/1/27
    no shutdown
    no routing
    vlan access 1
interface 1/1/28
    no shutdown
    no routing
    vlan access 1
interface vlan1
    ip dhcp
https-server vrf default
https-server vrf mgmt
vsf member 1
    type jl668a
6300#
```

3. Move to configuration mode and change the switch's hostname to **TX- Access-1**.

> **NOTE:** Replace the highlighted "X" in the hostname for your student table number, e.g. T4-Access-1 for table 4, or T11-Access-1 for table 11 like in the example below.

```
6300# configure terminal
6300(config)# hostname TX-Access-1
```

Access-1

4. Apply the console session timeout to 1 day (1440 minutes) to prevent a logout during the lab activities.

```
T11-Access-1(config)# session-timeout 1440
T11-Access-1(config)#
```

> **TIP:** An alternative method you can use is the next configuration script:
>
> ```
> T11-Access1(config)# cli-session
> ```

```
T11-Access1(config-cli-session)# timeout 1440
T11-Access1(config-cli-session)# exit
```

5. Use “**show interface brief”** for displaying a table of ports and their more relevant settings.

```
T11-Access-1# show interface brief
--------------------------------------------------------------------------------
----
Port       Native  Mode   Type           Enabled Status  Reason              Speed
           VLAN                                                             (Mb/s)
--------------------------------------------------------------------------------
1/1/1      1       access 1GbT           yes     up                           1000
1/1/2      1       access 1GbT           yes     down                         1000
1/1/3      1       access 1GbT           yes     up      Waiting for link       --
1/1/4      1       access 1GbT           yes     down    Waiting for link       --
                             ←---- output omitted ---→
1/1/24     1       access 1GbT           yes     down    Waiting for link       --
1/1/25     1       access SFP+DAC1       yes     up                          10000
1/1/26     1       access SFP+DAC1       yes     up                          10000
1/1/27     1       access SFP+DAC1       yes     up                          10000
1/1/28     1       access SFP+DAC1       yes     up                          10000
vlan1      --             --             yes     up                            --
```

What are the ports “Mode” values?

____

What ports are enabled?

____

NOTE: 6300 and 6400 AOS-CX switches have all their ports configured as layer 2 interfaces (VLAN and Spanning Tree capable) and are enabled by default vs 8000 switches that have administratively disabled routed ports.

6. Disable ports **1/1/2 to 1/1/28**.

```
T11-Access-1(config)# interface 1/1/2-1/1/28
T11-Access-1(config-if-<1/1/2-1/1/28>)# shutdown
```

```
T11-Access-1(config-if-<1/1/2-1/1/28>)# exit
```

7. Enable port **1/1/3**.

```
T11-Access-1(config)# interface 1/1/3
T11-Access-1(config-if)# no shutdown
T11-Access-1(config-if)# exit
```

8. Issue the "**show interface brief**" command again.

```
T11-Access-1(config)# show interface brief
-------------------------------------------------------------------------------
Port       Native  Mode   Type           Enabled Status  Reason           Speed
           VLAN                                                           (Mb/s)
-------------------------------------------------------------------------------
1/1/1      1       access 1GbT           yes     up                       1000
1/1/2      1       access 1GbT           no      down    Administratively down   --
1/1/3      1       access 1GbT           yes     up                       1000
1/1/4      1       access 1GbT           no      down    Administratively down   --
                             ←---- output omitted ---→
1/1/24     1       access 1GbT           no      down    Administratively down   --
1/1/25     1       access SFP+DAC1       no      down    Administratively down   --
1/1/26     1       access SFP+DAC1       no      down    Administratively down   --
1/1/27     1       access SFP+DAC1       no      down    Administratively down   --
1/1/28     1       access SFP+DAC1       no      down    Administratively down   --
vlan1      --             --             yes     up                              --
```

What is the Enabled and Status values of ports 1/1/27 and 1/1/28 now?

____

9. Display the "**event log**" in reverse mode.

```
T11-Access-1(config)# show events -r -n 10
---------------------------------------------------
Event logs from current boot
---------------------------------------------------
2020-01-08T20:23:22.656240+00:00 T11-Access-1 lldpd[2773]:
Event|106|LOG_INFO|MSTR|1|LLDP neighbor 88:3a:30:97:a4:40 deleted on 1/1/28
2020-01-08T20:23:22.656240+00:00 T11-Access-1 lldpd[2773]:
Event|106|LOG_INFO|MSTR|1|LLDP neighbor 88:3a:30:97:a4:40 deleted on 1/1/27
2020-01-08T20:23:11.652815+00:00 T11-Access-1 lldpd[2773]:
Event|106|LOG_INFO|MSTR|1|LLDP neighbor 90:20:c2:bc:26:00 deleted on 1/1/26
```

```
2020-01-08T20:23:32.660046+00:00 T11-Access-1 lldpd[2773]:
Event|106|LOG_INFO|MSTR|1|LLDP neighbor 90:20:c2:bc:ee:00 deleted on 1/1/27
2020-01-08T20:21:27.949674+00:00 T11-Access-1 intfd[715]:
Event|404|LOG_INFO|||Link status for interface 1/1/28 is down
2020-01-08T20:21:27.949674+00:00 T11-Access-1 intfd[715]:
Event|404|LOG_INFO|||Link status for interface 1/1/27 is down
2020-01-08T20:21:27.938854+00:00 T11-Access-1 intfd[715]:
Event|404|LOG_INFO|||Link status for interface 1/1/26 is down
2020-01-08T20:21:27.921704+00:00 T11-Access-1 intfd[715]:
Event|404|LOG_INFO|||Link status for interface 1/1/25 is down
2020-01-08T20:21:06.717715+00:00 T11-Access-1 hpe-restd[702]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-01-08T20:16:06.716627+00:00 6300 hpe-restd[702]: Event|4646|LOG_ERR|AMM|-
|Aruba Activate server https://devices-v2.arubanetworks.com is not reachable
through any supported VRF.
```

What link stats messages can you see at top related to 1/1/27 and 1/1/28 ports?

____

____

What other messages in the event log do you get?

____

____

____

ANSWER: You should see notifications informing you that LLDP neighbors have been deleted, because the ports have been disabled. Also, since AOS-CX switches periodically attempt to contact the Aruba Activate Cloud service and the switch has no internet connectivity the device complains that the service is unreachable.

10. Define interface descriptions for port **1/1/1** and **1/1/3**. Do not leave interface 1/1/3 yet.

```
T11-Access-1# configure terminal
```

```
T11-Access-1(config)# interface 1/1/1
T11-Access-1(config-if)# description TO_PC-1
T11-Access-1(config-if)# interface 1/1/3
T11-Access-1(config-if)# description TO_PC-3
T11-Access-1(config-if)#
```

11. Inside of interface 1/1/3 type the "**show running-config current-context**" command.

```
T11-Access-1(config-if)# show running-config current-context
interface 1/1/3
    no shutdown
    description TO_PC-3
    no routing
    vlan access 1
    exit
```

IMPORTANT: This command is a shortcut for displaying only the commands available at the context/subcontext level. Get used to it, since it is of great use when configuring and editing ports, protocols, access control lists, etcetera.

12. Run the "**show interface 1/1/3**" command followed by **"| include Description**".

NOTE: The information will be filtered out, listing the lines that include the "Description" string only, hence it is removing any other line part of that command's regular output.

```
T11-Access-1(config-if)# end
T11-Access-1# show interface 1/1/3 | include Description
 Description: TO_PC-3
```

NOTICE: The pipe (|) command filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.

Strings of characters that follow the filtering tool (e.g. “Description” in command above) are case sensitive. Typing the wrong capitalization may lead to the absence of output.

13. Try the same command but use **“| begin 3 Interface**” instead.

NOTE: The information will be filtered out, listing only the lines that include the “Interface” string along with the 3 subsequent lines.

```
T11-Access-1# show interface 1/1/3 | begin 3 Interface
Interface 1/1/3 is up
 Admin state is up
 Link transitions: 1
 Description: TO_PC-3
T11-Access-1(config-if)# end
```

How was the output modified now?

Task 3: Create and Explore Checkpoints.

Objectives

You have made some configuration changes in 6300-A, now is a good time to keep those changes stored in the system and protect them from any power cycle events. Next you will explore checkpoints, see how they are created, and make your own to save your progress.

Steps

Access-1

1. Open a console connection to **Access-1**.
2. Show the current system's checkpoints.

```
T11-Access-1# show checkpoint list
CPC20200108211347
CPC20200108212625
T11-Access-1#
```

How many entries did you get?

____

IMPORTANT: AOS-CX systems are 100% database driven. This means that configuration scripts you save are stored in a local database instead of a regular configuration file. The database is periodically tracked and whenever the changes are made, they will be automatically stored after a 5-minute idle period. Any new configuration change, followed by a 5-minute idle period, will create a new checkpoint that can later be used to back up or restore the running configuration state of the system.

On demand checkpoints can be generated by saving the running-configuration or creating custom checkpoints.

3. Issue the “**write memory**” command.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

4. List the checkpoints again.

```
T11-Access-1# show checkpoint list
CPC20200108211347
CPC20200108212625
startup-config
T11-Access-1#
```

Is there any new checkpoint?

____

What is its name?

____

5. Create a checkpoint called **Lab3** using the **running-configuration** as the source.

```
T11-Access-1# copy running-config checkpoint Lab3
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

6. Display the checkpoints one more time.

```
T11-Access-1# show checkpoint list
CPC20200108211347
CPC20200108212625
startup-config
Lab3
```

7. Now make a checkpoint called **Lab3_final** using the **running-config** as the source.

```
T11-Access-1# copy running-config checkpoint Lab3_final
Configuration changes will take time to process, please be patient.
cannot create duplicate checkpoint, configuration already exists in checkpoint
Lab3
T11-Access-1#
```

What error message did you get?

____

NOTE: AOS-CX cannot have two different configuration snapshots with identical contents in its database (that would not be resource efficient). If you want to rename a checkpoint, then you will have to delete it first, then create a new one.

8. Erase checkpoint Lab3.

```
T11-Access-1# erase checkpoint Lab3
```

9. Try creating the checkpoint again.

```
T11-Access-1# copy running-config checkpoint Lab3_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

10. Last issue the “**show checkpoint list all**” command.

```
T11-Access-1# show checkpoint list all
|NAME               |TYPE       |WRITER  |DATE(UTC)              …|IMAGE VERSION|
|CPC20200108211347  |checkpoint |System  |2020-01-08T21:13:47Z   …|FL.10.04.0030|
|CPC20200108212625  |checkpoint |System  |2020-01-08T21:26:25Z   …|FL.10.04.0030|
|startup-config     |startup    |User    |2020-01-08T22:18:31Z   …|FL.10.04.0030|
|Lab3_final         |latest     |User    |2020-01-08T22:28:33Z   …|FL.10.04.0030|
```

IMPORTANT:

You will see the same list of checkpoints along with more detailed data about them, like checkpoint type, user who created it, date and time it was created and OS release that was running when they were created. Keeping track of when checkpoints are created is important during regular maintenance tasks. This is the reason configuring all switches with Network Time Protocol server is important.

Since IP connectivity is not enabled yet, you will continue working without setting up an NTP server and trust the system clock for now. NTP configuration will be covered in a later Module.

IMPORTANT:

Checkpoints can be restored by using the copy command and applying the checkpoint's contents into the running-configuration (or startup configuration and invoking the "boot system" command), like in the example below.

```
T11-Access-1# copy checkpoint Lab3_final running-config
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

You have completed Lab 3!

AOS-CX Switching Fundamentals

Lab 4.1: Configure a VLAN

Overview

At this point the Access-1 switch is up and running and ready for configuration. The next task in your initial network deployment will be to place wired employees in a custom VLAN in order to enable inter user communication.

Objectives

After completing this lab, you will be able to:

- Create a custom VLAN and assign it to access ports
- Configure clients with static IP addresses
- Explore the Switch MAC address table

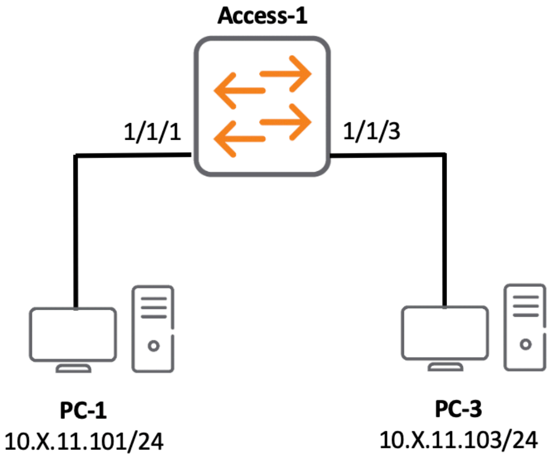

Figure 4.1-1: Lab Topology

Task 1: Assign PCs to VLAN X11

Objectives

In this task you will create the employee VLAN and configure Windows PCs with IP addresses of the corresponding IP segment according to the network design. Then you will verify IP connectivity between clients and explore the MAC address table.

Steps

Access-1

1. Open a console connection to **Access-1.** Login with **admin** and **no password**.
2. Use the "**show vlan**" command to display current Virtual Local Area Networks configured in the switch. You should only see VLAN 1 assigned to all ports. This is the default setting for the switch.

```
T11-Access-1# show vlan

-----------------------------------------------------------------------------------
VLAN  Name                        Status  Reason                 Type     Interfaces
-----------------------------------------------------------------------------------
1     DEFAULT_VLAN_1              up      ok                     default  1/1/1-1/1/28
T11-Access-1#
```

3. Create **VLAN X11** and name it **EMPLOYEES**.

> **NOTE:** Replace the highlighted "X" in the vlan command for your student table number, e.g. "vlan 411" for table 4, or "vlan 1111" for table 11 like in the example below.

```
T11-Access-1# configure terminal
T11-Access-1(config)# vlan X11
T11-Access-1(config-vlan-1111)# name EMPLOYEES
T11-Access-1(config-vlan-1111)# exit
```

4. Repeat the show vlan command.

```
T11-Access-1(config)# show vlan

---------------------------------------------------------------------------------
VLAN  Name                    Status  Reason                Type     Interfaces
---------------------------------------------------------------------------------
1     DEFAULT_VLAN_1          up      ok                    default  1/1/1-1/1/28
1111  EMPLOYEES               down    no_member_port  static
T11-Access-1(config)#
```

Is the output reflecting your previous configuration change?

____

What is the newly created VLAN status?

____

What caused the new VLAN to have this status?

____

ANSWER: Since the VLAN has not been assigned to any enabled physical port, the status is down. No MAC address learning process is happening in the switch for that VLAN.

5. Assign **VLAN X11** to interfaces **1/1/1 and 1/1/3** as an access VLAN.

```
T11-Access-1(config)# interface 1/1/1
T11-Access-1(config-if)# vlan access X11
T11-Access-1(config-if)# interface 1/1/3
T11-Access-1(config-if)# vlan access X11
T11-Access-1(config-if)# exit
```

6. Try the “**show vlan**” command again.

```
T11-Access-1(config)# show vlan

---------------------------------------------------------------------------------
```

```
VLAN  Name                    Status  Reason                Type     Interfaces
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1          down    no_member_forwarding  default  1/1/2,1/1/4-
1/1/28
1111  EMPLOYEES               up      ok                    static   1/1/1,1/1/3
T11-Access-1(config)#
```

What is the VLAN X11 status now?

____

NOTE: Currently, only ports 1/1/1 and 1/1/3 are UP. When you replaced VLAN 1 with VLAN X11 on the ports, both VLANs will still appear, but VLAN 1 is no longer associated with any port in the UP state. Therefore, VLAN 1's status was changed to down.

7. Issue the "**show vlan port 1/1/1**" command.

```
T11-Access-1(config)# show vlan port 1/1/1

------------------------------------------------------------------------------
VLAN  Name                            Mode            Mapping
------------------------------------------------------------------------------
1111  EMPLOYEES                       access          port
T11-Access-1(config)#
```

What VLAN is present on the interface and what is its mode?

____

8. Use the "**show vlan summary**" command. This command shows the VLAN count in the system.

```
T11-Access-1(config)# show vlan summary
Number of existing VLANs: 2
Number of static VLANs:   2
Number of dynamic VLANs:  0
```

9. Issue the "**show interface 1/1/1**" command. You will be able to see VLAN ID and VLAN Mode at the bottom of the command.

```
T11-Access-1(config)# show interface 1/1/1

Interface 1/1/1 is up
 Admin state is up
 Link transitions: 1t
 Description: TO_PC-1
 Hardware: Ethernet, MAC Address: 88:3a:30:98:30:27
 MTU 1500
 Type 1GbT
 Full-duplex
 qos trust none
 Speed 1000 Mb/s
 Auto-negotiation is on
 Flow-control: off
 Error-control: off
 MDI mode: MDIX
 VLAN Mode: access
 Access VLAN: 1111
```

10. Finally, try the "**show interface brief**" command followed by a filtering option "**| begin 5 Port**".

> **NOTE:** The information will be filtered out, listing only the lines that include the "Port" string along with the 5 subsequent lines.

```
T11-Access-1(config)# show interface brief | begin 5 Port
Port      Native  Mode   Type           Enabled Status  Reason              Speed
          VLAN
(Mb/s)
--------------------------------------------------------------------------------
1/1/1     1111    access 1GbT           yes     up                           1000
1/1/2     1       access 1GbT           no      down    Administratively down  --
1/1/3     1111    access 1GbT           yes     up                           1000
T11-Access-1(config)#
```

> **NOTE:** The pipe (|) command filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.
>
> Strings of characters that follow the filtering tool (e.g. "Port" in the example above) are case sensitive. Incorrect capitalization may lead to the absence of output or other unexpected result.

What is the value under Native VLAN for ports 1/1/1 and 1/1/3 vs 1/1/2?

Task 2: Explore MAC Address Table

Objectives

In this second task, you will statically define IP addresses to PC-1 and PC-2, so they can achieve intra VLAN layer 3 connectivity, and users on those machines can start collaborating to run their company's daily operations.

Steps

PC-1

1. Access **PC-1**'s console.
2. Under search field in the task bar, type "**control panel"**. Windows will automatically display all items matching the string.
3. Click the top result (**Control Panel**). A new window will pop up.

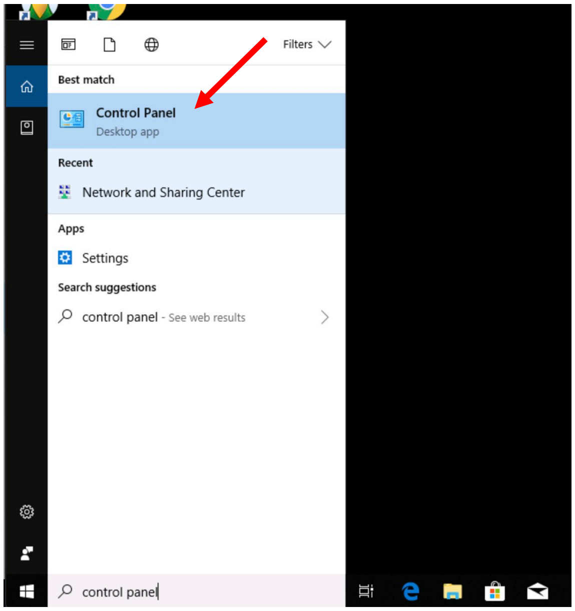

Figure 4.1-2: Windows Search

4. In Control Panel, click "**View network status and tasks**" under Network and Internet.

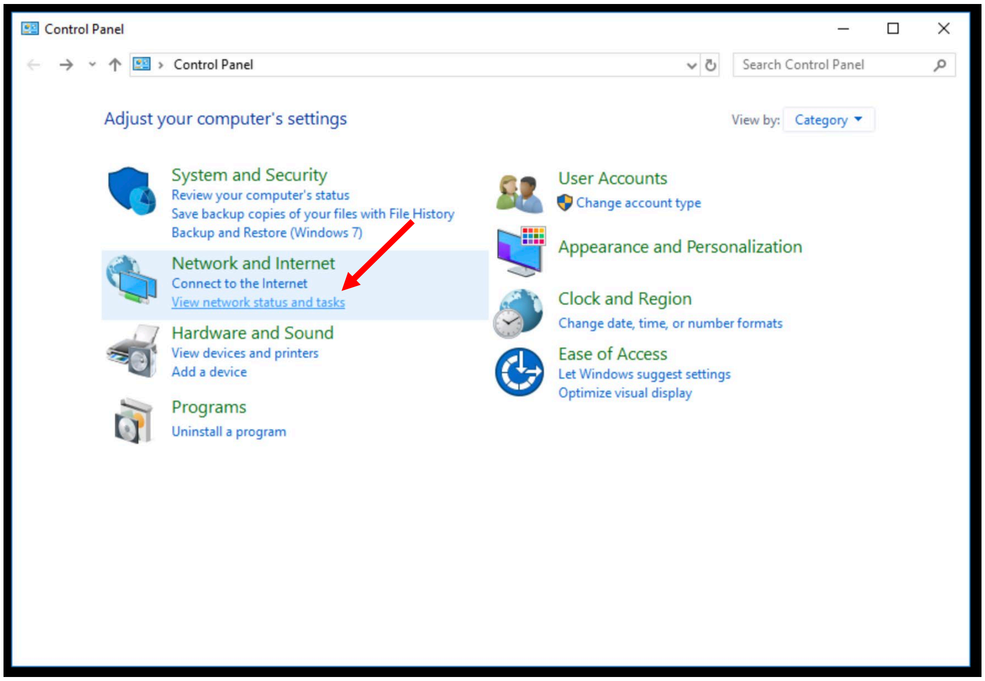

Control Panel

Control Panel

Search Control Panel

Adjust your computer's settings

View by: Category

System and Security
Review your computer's status
Save backup copies of your files with File History
Backup and Restore (Windows 7)

Network and Internet
Connect to the Internet
View network status and tasks

Hardware and Sound
View devices and printers
Add a device

Programs
Uninstall a program

User Accounts
Change account type

Appearance and Personalization

Clock and Region
Change date, time, or number formats

Ease of Access
Let Windows suggest settings
Optimize visual display

Figure 4.1-3: Windows Control Panel

5. Click at "**Change adapter settings**" on the left options.

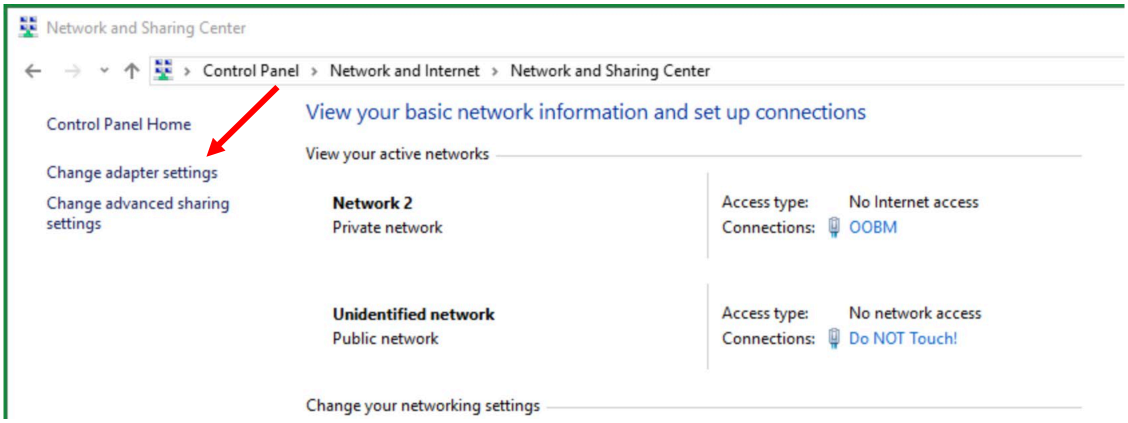

Network and Sharing Center

Control Panel › Network and Internet › Network and Sharing Center

Control Panel Home

Change adapter settings

Change advanced sharing settings

View your basic network information and set up connections

View your active networks

Network 2
Private network

Access type: No Internet access
Connections: OOBM

Unidentified network
Public network

Access type: No network access
Connections: Do NOT Touch!

Change your networking settings

Figure 4.1-4: Change adapter settings

6. Double click "**Lab NIC**" to access the NIC Status Window.

NOTE: If NIC was disabled (greyed out), then you will have to double click it twice, first to enable it then a second time to access the NIC Status Window.

NOTICE: There is an interface called "Do NOT Touch!", please repeat with me, "do not touch!!!" If changes are made to that NIC (like modifying the IP address or disabling the interface) the access to this virtual machine will be disrupted. Only the lab support team will be able to recover the system and that process may delay your lab progress.

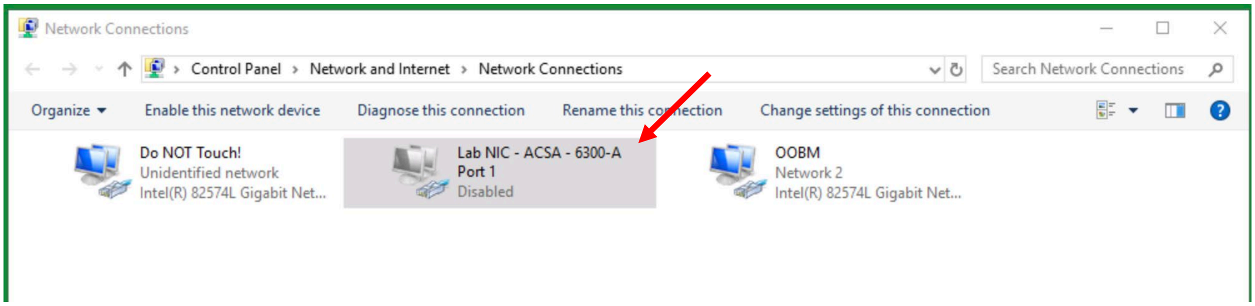
Network Connections
Control Panel > Network and Internet > Network Connections
Search Network Connections
Organize
Enable this network device
Diagnose this connection
Rename this connection
Change settings of this connection
Do NOT Touch!
Unidentified network
Intel(R) 82574L Gigabit Net...
Lab NIC - ACSA - 6300-A Port 1
Disabled
OOBM
Network 2
Intel(R) 82574L Gigabit Net...

Figure 4.1-5: Network and Sharing Center

7. In Lab NIC status window, click "**Properties**" button.

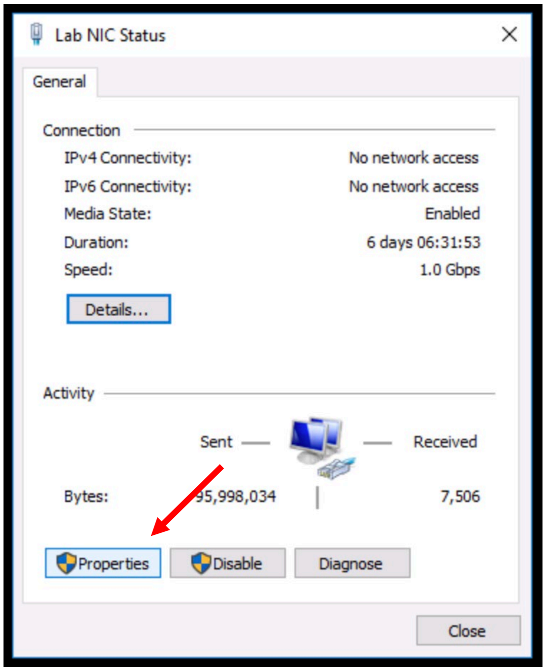

Lab NIC Status

General

Connection

IPv4 Connectivity:	No network access
IPv6 Connectivity:	No network access
Media State:	Enabled
Duration:	6 days 06:31:53
Speed:	1.0 Gbps

Figure 4.1-6: Lab NIC Status

8. In Lab NIC Properties section, select “**Internet Protocol Version 4 (TCP/IPv4)**, then click “**Properties**” button.

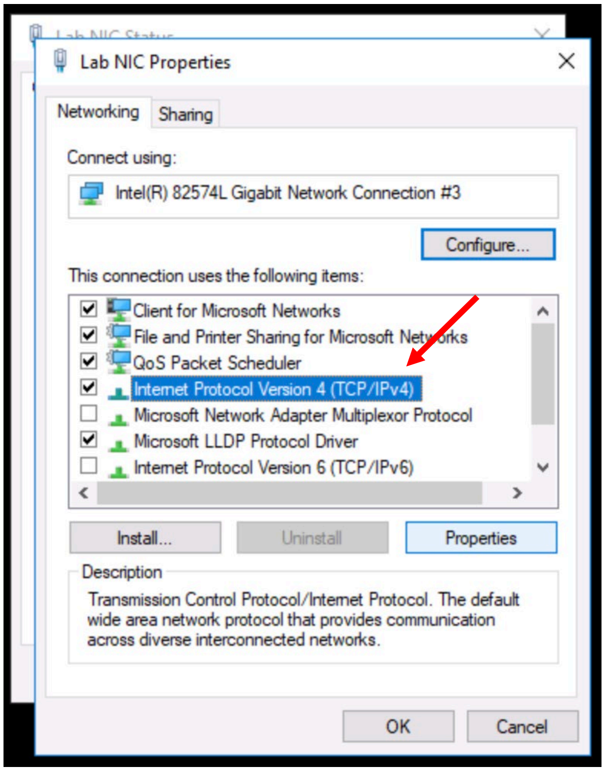

Lab NIC Properties

Connect using:

Intel(R) 82574L Gigabit Network Connection #3

This connection uses the following items:

☑ Client for Microsoft Networks
☑ File and Printer Sharing for Microsoft Networks
☑ QoS Packet Scheduler
☑ Internet Protocol Version 4 (TCP/IPv4)
□ Microsoft Network Adapter Multiplexor Protocol
☑ Microsoft LLDP Protocol Driver
□ Internet Protocol Version 6 (TCP/IPv6)

Description

Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

Figure 4.1-7: Lab NIC Properties

9. In Internet Protocol Version 4 (TCP/IPv4) Properties, choose “**Use the following IP address:**” under General tab.

10. Type **10.X.11.101** and **255.255.255.0** under IP address and Subnet mask respectively.

NOTE: Replace the highlighted “X” with your student table number, e.g. 10.4.11.101 for table 4, or 10.11.11.101 for table 11 like in the example below.

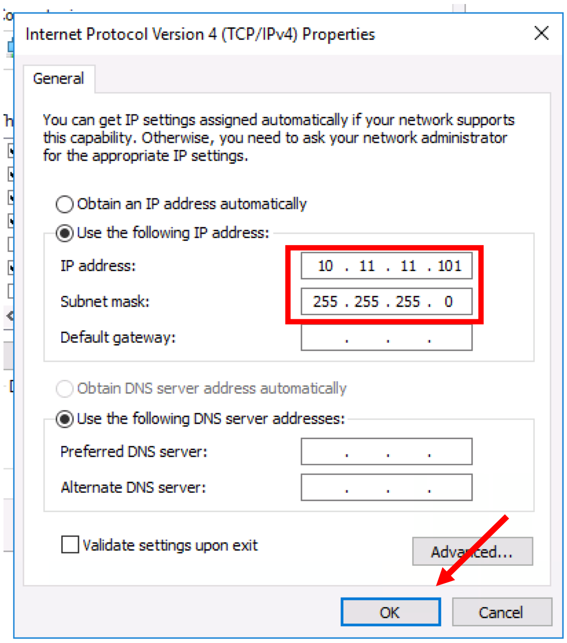

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 11 . 101

Subnet mask: 255 . 255 . 255 . 0

Default gateway: . . .

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 4.1-8: Internet Protocol Version 4 Properties

11. Click “**OK**” button, then “**Close”** button twice.
12. Under search field in the task bar, type “**command”**. Windows will automatically display all items matching the string.

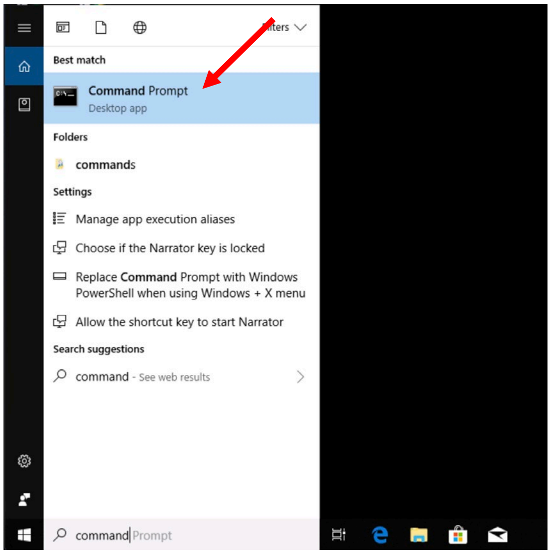

Best match

Command Prompt
Desktop app

Folders

commands

Settings

Manage app execution aliases

Choose if the Narrator key is locked

Replace **Command** Prompt with Windows PowerShell when using Windows + X menu

Allow the shortcut key to start Narrator

Search suggestions

command - See web results

command Prompt

Figure 4.1-9: Windows Search 2

13. Click the top result (**Command Prompt**). A new window will pop up.

14. Type “**ipconfig**” and hit **[Enter].** This command will display IPv4 settings of all NICs in the system.
15. Confirm the Ethernet adapter called Lab NIC has the IPv4 address you just configured.

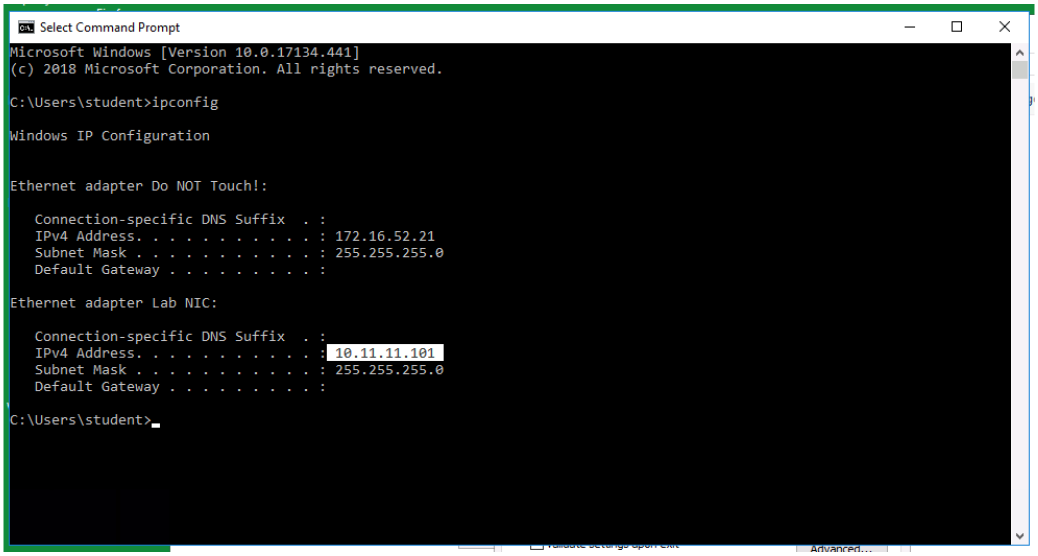

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ipconfig

Windows IP Configuration


Ethernet adapter Do NOT Touch!:

   Connection-specific DNS Suffix  . :
   IPv4 Address. . . . . . . . . . . : 172.16.52.21
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :

Ethernet adapter Lab NIC:

   Connection-specific DNS Suffix  . :
   IPv4 Address. . . . . . . . . . . : 10.11.11.101
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :

C:\Users\student>_
```

Figure 4.1-10: ipconfig

16. Type “**ipconfig -all**” version of the command and hit **[Enter].** This command displays additional information like DNS servers IP addresses (if configured) and the NICs physical address (MAC).

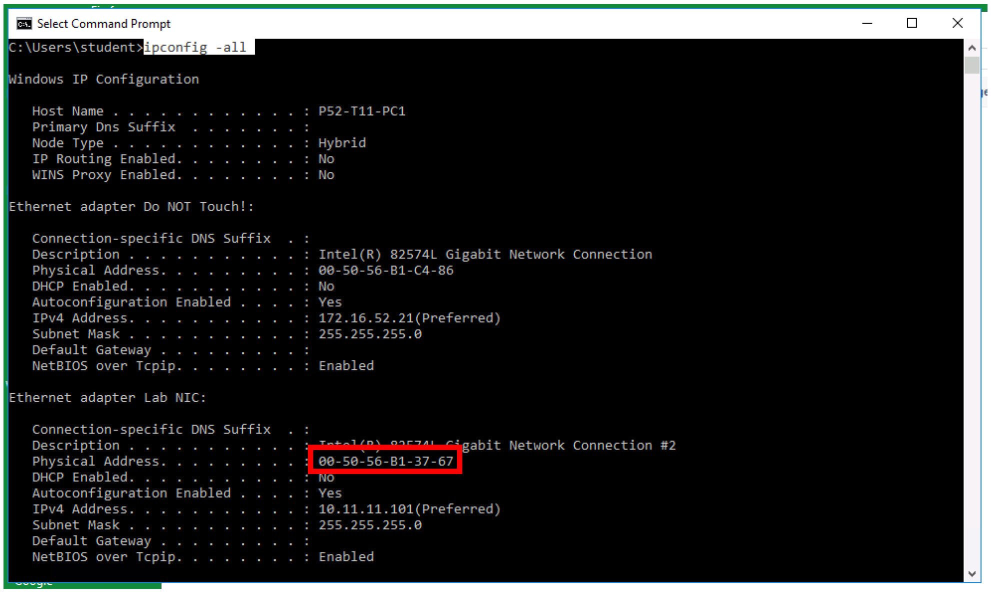

```
C:\Users\student>ipconfig -all

Windows IP Configuration

   Host Name . . . . . . . . . . . . : P52-T11-PC1
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No

Ethernet adapter Do NOT Touch!:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection
   Physical Address. . . . . . . . . : 00-50-56-B1-C4-86
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 172.16.52.21(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   NetBIOS over Tcpip. . . . . . . . : Enabled

Ethernet adapter Lab NIC:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection #2
   Physical Address. . . . . . . . . : 00-50-56-B1-37-67
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 10.11.11.101(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

Figure 4.1-11: ipconfig -all

What is PC-1's Lab NIC MAC address?

____

This is the typical IP address configuration process in a Windows system. You will now repeat it on PC-3

PC-3

17. Access **PC-3**'s console and repeat **steps 2 to 11** using **10.X.11.103** IP address instead.
18. If there is any "**OOBM**" NIC, then disable it.

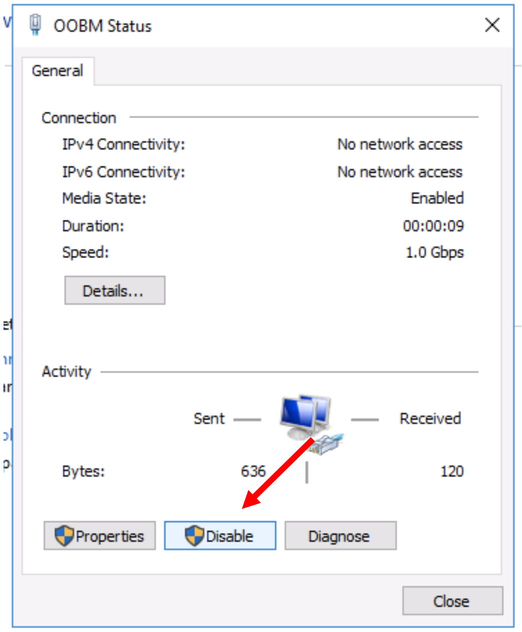

OOBM Status

General

Connection

IPv4 Connectivity: No network access

IPv6 Connectivity: No network access

Media State: Enabled

Duration: 00:00:09

Speed: 1.0 Gbps

Activity

Bytes: Sent 636 Received 120

Figure 4.1-12: NIC disabled

19. Repeat **steps 12 to 16**.

What is PC-3's Lab NIC MAC address?

20. Confirm “**OOBM**” NIC is not listed.
21. From PC-3, **ping PC-1**’s IP address (**10.X.11.101**). Ping should be successful.

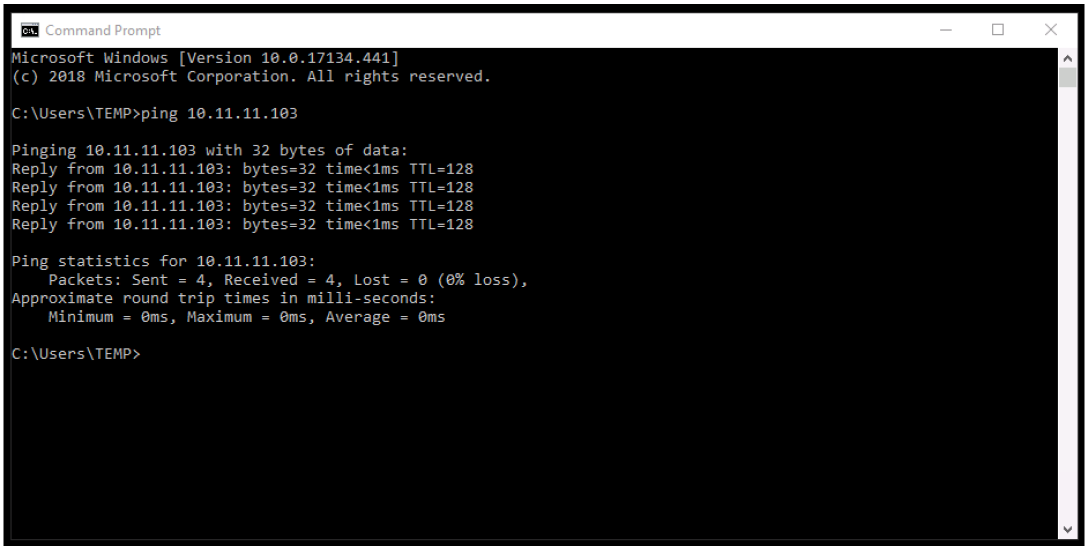

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\TEMP>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\TEMP>
```

Figure 4.1-13: Ping to PC-3

Access-1

22. In Access-1, display the mac-address-table.

```
T11-Access-1# show mac-address-table
MAC age-time                : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                      Port
--------------------------------------------------------------
00:50:56:b1:37:67    1111     dynamic                   1/1/1
00:50:56:b1:ae:e8    1111     dynamic                   1/1/3
T11-Access-1#
```

What entries are listed in the output?

23. Using the output information, write down the client’s MAC addresses in figure 4.1-14, along with ports and VLAN IDs.

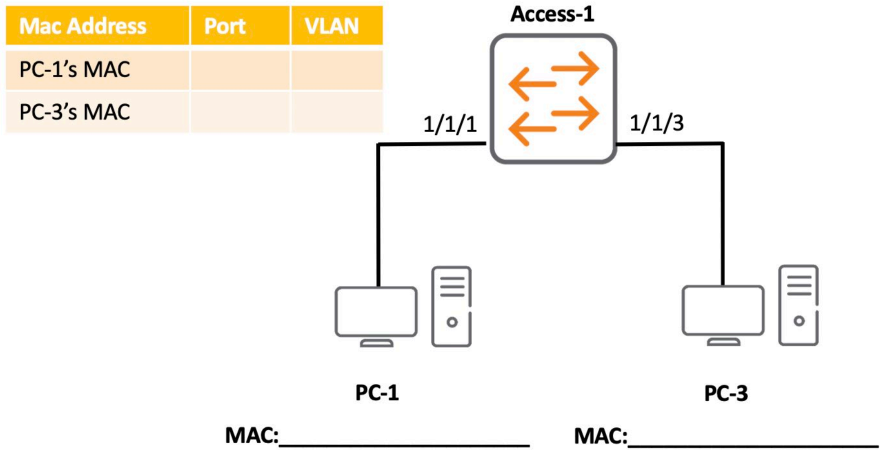

Mac Address	Port	VLAN
PC-1’s MAC		
PC-3’s MAC		

Figure 4.1-14: Access-1’s MAC Address Table

TIP: You can find a larger copy of this diagram in Appendix 3.

Were these MAC addresses discovered on the ports you expected them?

____

TIP: There are multiple forms of the “show mac-address-table” command that can be used for displaying only entries that match a certain criteria, such as an address learned in a particular VLAN or port, or learned dynamically versus configured statically in the MAC table, use the [?] key at the end of the command for displaying the options.

```
T11-Access-1# show mac-address-table ?
     address   Show a specific MAC address
     count     Number of MAC addresses
     detail
     dynamic   Show learnt MAC addresses
     lockout   Show MAC lockout address information
     port      Show MAC addresses learnt on port
     static    Show static MAC address information
     vlan      Show MAC addresses learnt on VLANs
     <cr>
```

Task 3: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Please note that final lab checkpoints may be used in later activities.

Steps

Access-1

1. Save the current Access-1's configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

2. Backup the current Access-1's configuration as a custom checkpoint called **Lab4-1_final**.

```
T11-Access-1# copy running-config checkpoint Lab4-1_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

You have completed Lab 4.1!

AOS-CX Switching Fundamentals

Lab 4.2: Add a Second Switch to the Topology

Overview

Good news! BigStartup seems to be a successful business and management has decided to hire more personnel. More ports are required and it is time to add a second switch. You have been asked to make an onsite visit to integrate the second switch and span the employee VLAN.

Objectives

After completing this lab, you will be able to:

- Enable an Interswitch link
- Configure trunk ports by enabling 802.1Q tagging on them
- Extend the broadcast domain
- Enable Inter-switch client communication

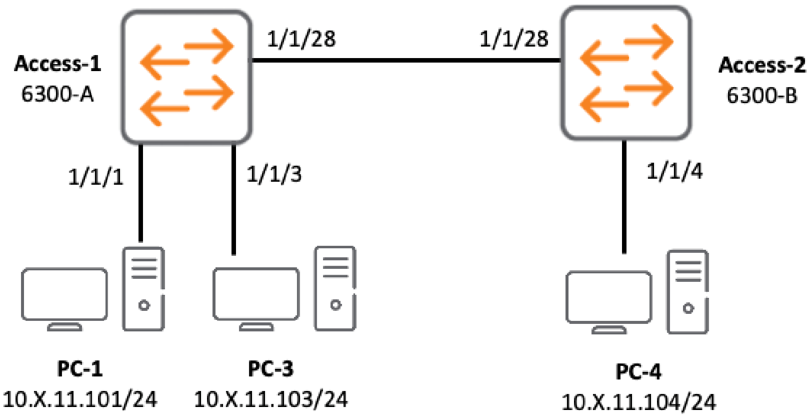

Figure 4.2-1: Lab Topology

Task 1: Configure Initial Settings on TX-Access-2

Objectives

Task 1 of lab 4.2 defines the initial settings for Access-2 and disables all ports but the one for the Windows client. Then you will move to PC-4 and assign an IP address to its NIC.

Steps

6300-B

1. Open a console connection to the **6300-B**. Login using **admin** and **no password**.
2. Move to configuration mode and change the switch's hostname to **TX-Access-2** and set session timeout to 1440 minutes.

> **NOTE:** Replace the highlighted "X" in the hostname for your student table number, e.g. T4- Access-1 for table 4, or T11-Access- for table 11 like in the example below.

```
6300# configure terminal
6300(config)# hostname TX-Access-2
T11-Access-2(config)# session-timeout 1440
T11-Access-2(config)#
```

3. Disable all ports.

```
T11-Access-2(config)# interface 1/1/1-1/1/28
T11-Access-2(config-if-<1/1/2-1/1/28>)# shutdown
T11-Access-2(config-if-<1/1/2-1/1/28>)# exit
```

4. Access interface **1/1/4** and set a description (this interface connects to PC-4).

```
T11-Access-2(config)# interface 1/1/4
T11-Access-2(config-if)# description TO_PC-4
```

5. Enable the port.

```
T11-Access-2(config-if)# no shutdown
T11-Access-2(config-if)# exit
```

You will now give PC4 an IP address.

PC-4

6. Open a console to **PC-4.**
7. Click the top result (**Control Panel**). A new window will pop up.

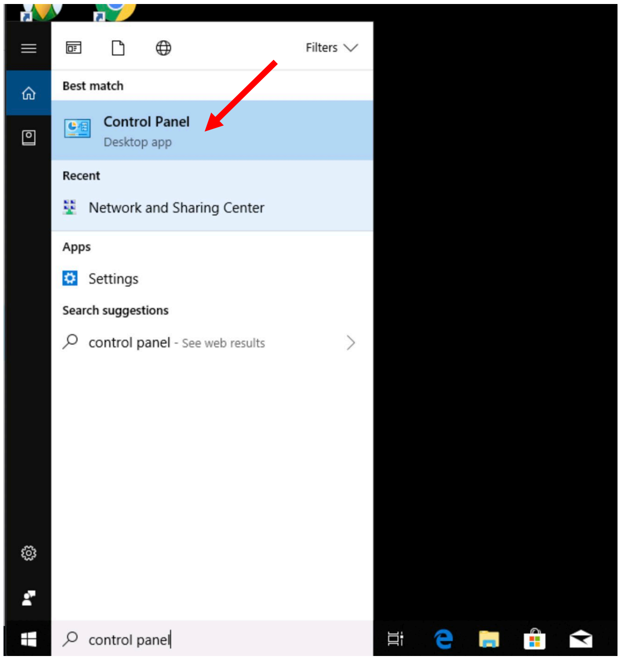

Figure 4.2-2: Windows search

8. Under Control Panel, click “**View network status and tasks**” under Network and Internet.

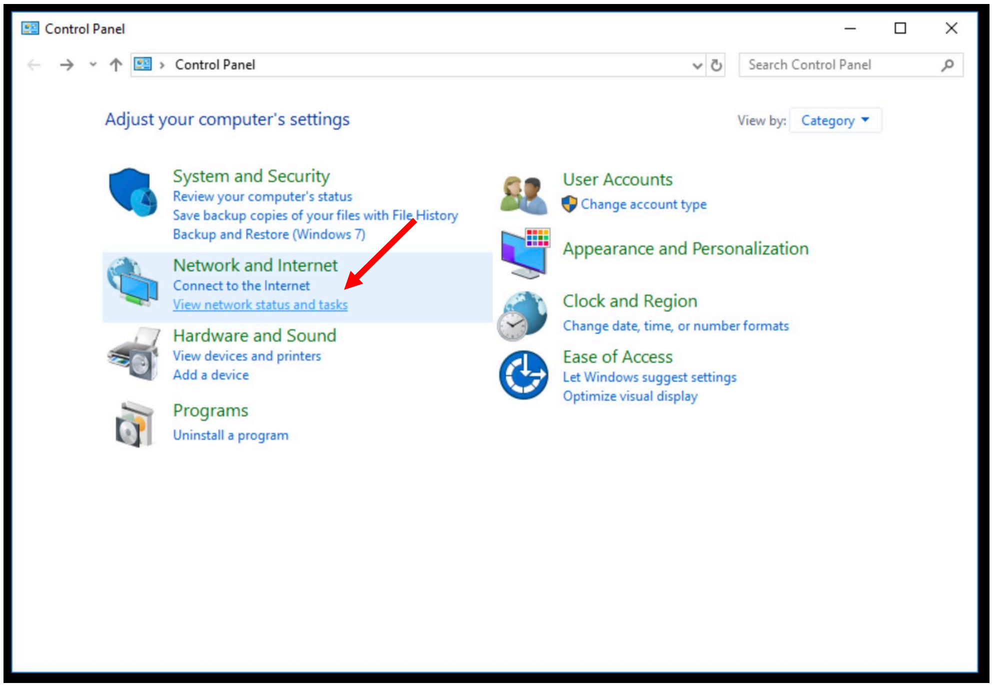
Control Panel
Control Panel
Search Control Panel
Adjust your computer's settings
View by: Category
System and Security
Review your computer's status
Save backup copies of your files with File History
Backup and Restore (Windows 7)
Network and Internet
Connect to the Internet
View network status and tasks
Hardware and Sound
View devices and printers
Add a device
Programs
Uninstall a program
User Accounts
Change account type
Appearance and Personalization
Clock and Region
Change date, time, or number formats
Ease of Access
Let Windows suggest settings
Optimize visual display

Figure 4.2-3: Control Panel

9. In Control Panel, click “**View network status and tasks**” under Network and Internet.
10. Click in “**Change adapter settings**” in the left pane.
11. Right click the “**Lab NIC**” adapter icon and select “**Properties**” from the menu that appears.

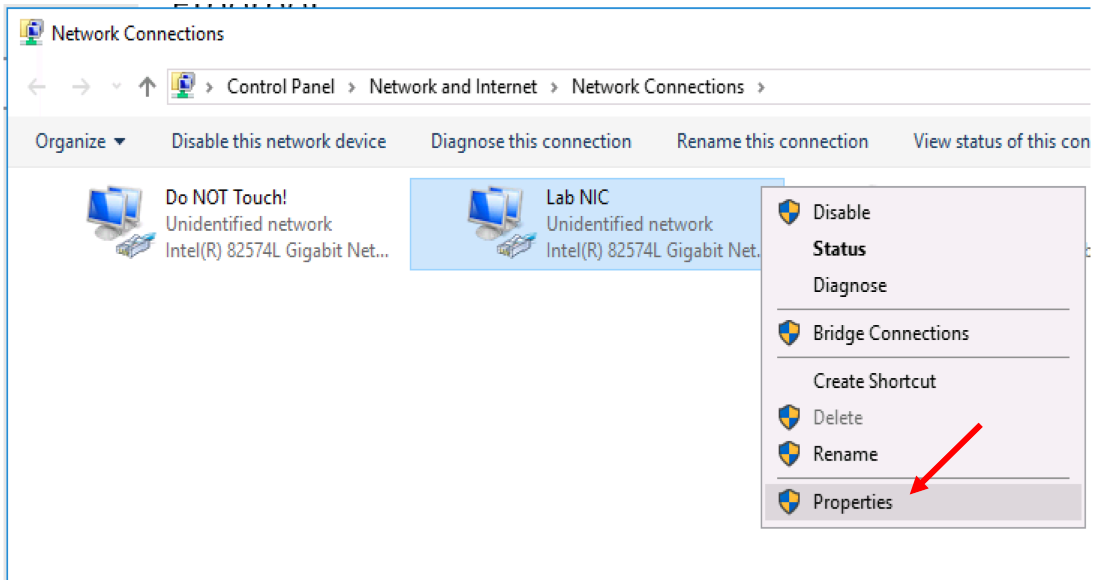

Figure 4.2-4: NIC Properties.

12. In Lab NIC status window, click “**Properties**” button.

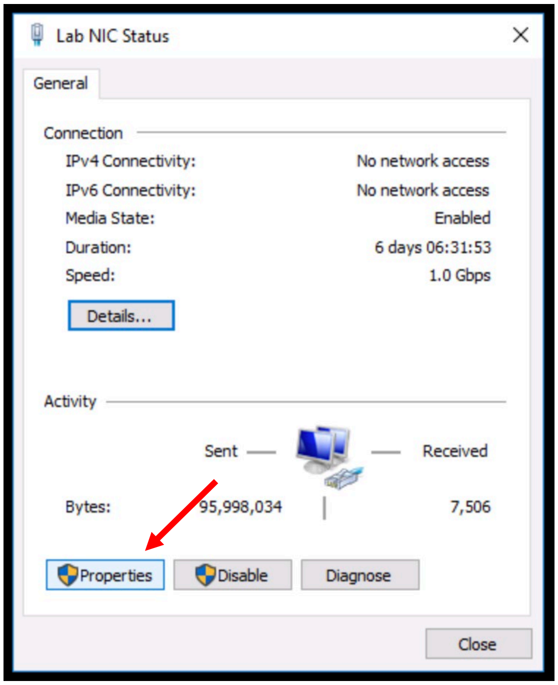

Lab NIC Status

General

Connection

IPv4 Connectivity: No network access

IPv6 Connectivity: No network access

Media State: Enabled

Duration: 6 days 06:31:53

Speed: 1.0 Gbps

Activity

Sent — Received

Bytes: 95,998,034 | 7,506

Figure 4.2-5: Lab NIC Status.

13. In Lab NIC Properties section, select “Internet **Protocol Version 4 (TCP/IPv4)**, then click “**Properties**” button.

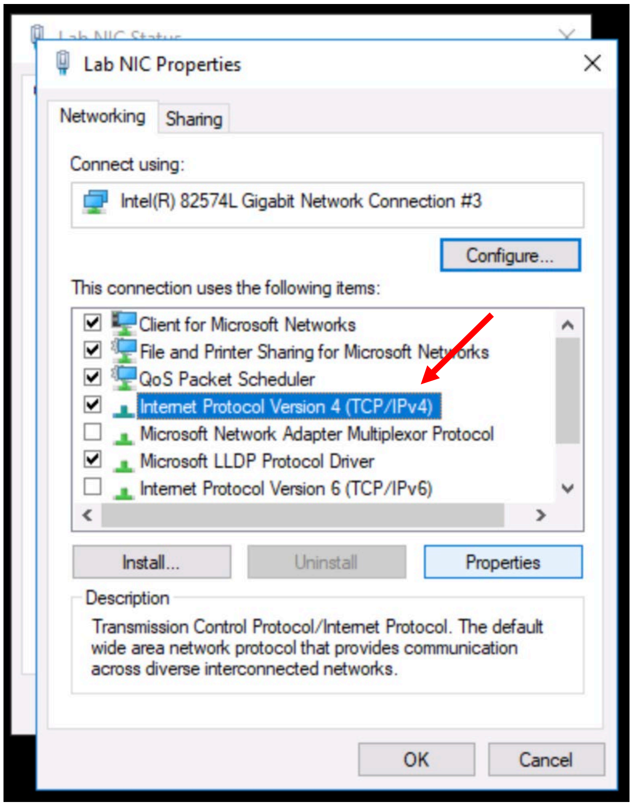

Lab NIC Properties

Networking | Sharing

Connect using:

Intel(R) 82574L Gigabit Network Connection #3

Configure...

This connection uses the following items:

☑ Client for Microsoft Networks
☑ File and Printer Sharing for Microsoft Networks
☑ QoS Packet Scheduler
☑ Internet Protocol Version 4 (TCP/IPv4)
☐ Microsoft Network Adapter Multiplexor Protocol
☑ Microsoft LLDP Protocol Driver
☐ Internet Protocol Version 6 (TCP/IPv6)

Install... | Uninstall | Properties

Description

Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

OK | Cancel

Figure 4.2-6: Lab NIC Properties.

14. In Internet Protocol Version 4 (TCP/IPv4) Properties, choose “**Use the following IP address:**” under General tab.
15. Type **10.X.11.104** and **255.255.255.0** under IP address and Subnet mask respectively.

NOTE: Replace “X” for your student table number, e.g. 10.4.11.104 for table 4, or 10.11.11.104 for table 11 like in the example below.

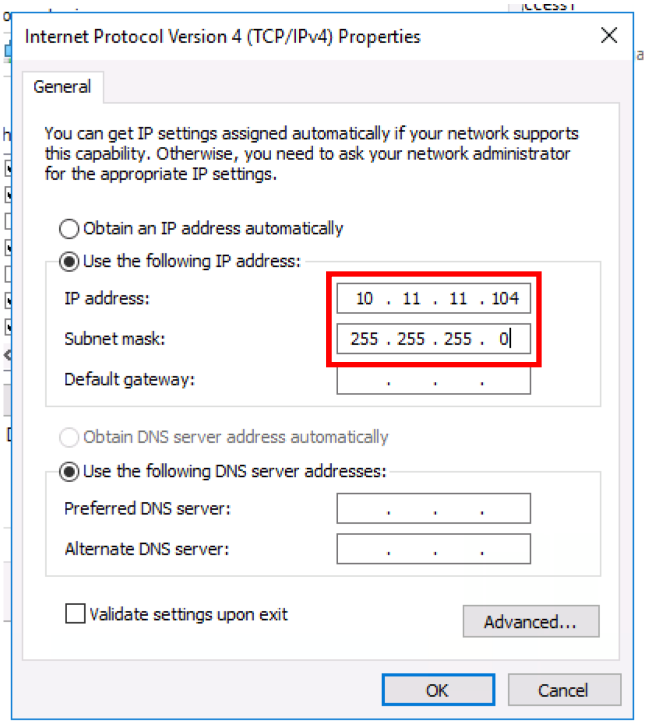

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 11 . 104

Subnet mask: 255 . 255 . 255 . 0

Default gateway: . . .

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 4.2-7: Internet Protocol Version 4.

16. Click "**OK**" button, then "**Close"** button twice.
17. Open the **Command Prompt**.
18. Ping **PC-3**'s IP address (**10.X.11.103**).

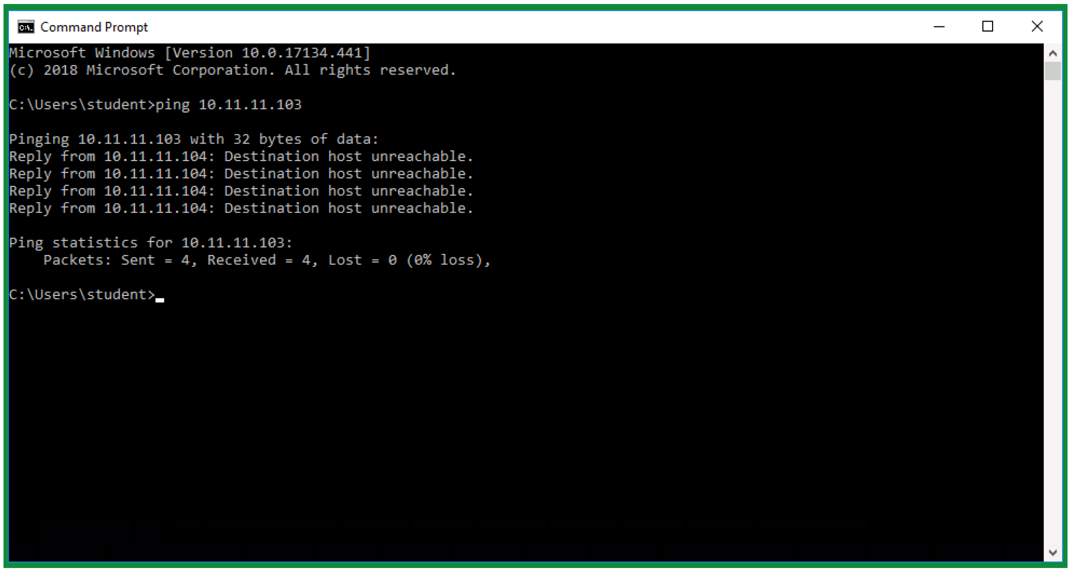

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\student>_
```

Figure 4.2-8: Ping failure.

NOTE: When destination IP address is within the source's IP segment and ping test result is "Destination host unreachable" it means that the Layer 3 to Layer 2 address resolution using Address Resolution Protocol (ARP) has failed and the ICMP echo message was not sent at all. However, if result is "timeout" then it means that host was able to resolve destination's MAC and ICMP packet was sent, but there is no reply coming back.

Was ping successful?

____

Why?

____

ANSWER: Ping is not successful because the destination IP address belongs to a device that is physically plugged into another switch (Access-1). Access-1 and Access-2 are not currently connected. Provisioning the Interswitch link in the next task will fix this issue.

Task 2: Enable Link Between Access Switches

In this task you will enable an ethernet connection between Access switches using a DAC in order to increase the number of ports on the network. Next, you will explore the information that Link Layer Discovery Protocol (LLDP) can provide.

Objectives

- Deploy a switch to switch link.
- Discover LLDP neighbors and look at detailed neighbor's information.
- Explore the switches' MAC Address tables.

Steps

Access-1

1. Open a console connection to the **Access-1.**
2. Enable interface 1/1/28.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/28
T11-Access-1(config-if)#no shutdown
T11-Access-1(config-if)#end
T11-Access-1(config-if)#
```

Access-2

3. Move to the **Access-2.**
4. Enable interface **1/1/28**.

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface 1/1/28
T11-Access-2(config-if)#no shutdown
T11-Access-2(config-if)#end
T11-Access-2#
```

5. Confirm interface 1/1/28 came up. Using the "**show interface brief**" command followed by the filter "**| exclude down**".

NOTE: The information will be filtered out, listing all the lines except the ones that contain the "down" string.

```
T11-Access-2# show interface brief | exclude down
--------------------------------------------------------------------------------
----
Port       Native  Mode   Type           Enabled Status  Reason             Speed
           VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
1/1/4      1112    access 1GbT           yes     up                         1000
1/1/28     1       access  SFP+DAC1      yes     up                         10000
vlan1      --             --             yes     up                         --
T11-Access-2#
```

NOTE: The pipe (|) command filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.

Strings of characters that follow the filtering tool (e.g. "down" in command above) are case sensitive. Typing the wrong capitalization may lead to the absence of output.

Is port 1/1/28 up?

____

What are port 1/1/4 and port 1/1/28 speeds?

____

IMPORTANT: In wired networking it is common practice to use faster speed links for connections between switches than those to the clients. Best practice for switch to switch connections is to limit oversubscription ratios to 24:1 or less (depending on the traffic generated by the endpoints). This guarantees that regardless of the traffic pattern, the link between switches does not get congested.

Next you will use LLDP to analyze the information the protocol can provide regarding what device is connected to specific interfaces.

NOTE: LLDP is on by default on AOS-CX switches.

6. Issue the "**show lldp configuration**" command.

```
T11-Access-2# show lldp configuration

LLDP Global Configuration
=========================

LLDP Enabled                 : Yes
LLDP Transmit Interval       : 30
LLDP Hold Time Multiplier    : 4
LLDP Transmit Delay Interval : 2
LLDP Reinit Time Interval    : 2

TLVs Advertised
===============

Management Address
Port Description
Port VLAN-ID
System Capabilities
System Description
System Name
OUI

LLDP Port Configuration
=======================

PORT           TX-ENABLED          RX-ENABLED
-----------------------------------------------
1/1/1          Yes                 Yes
1/1/2          Yes                 Yes
                              ←---- output omitted ---→
1/1/27         Yes                Yes
1/1/28         Yes                Yes
```

What is the current LLDP state?

____

What are the transmit interval and hold time multiplier values?

____

What are the LLDP transmit and receive modes on all of the ports?

____

NOTE: LLDP is enabled by default both globally and per port (on all ports). It can be disabled/enabled globally and/or per port using the commands shown below:

```
T11-Access-1(config)# no lldp
T11-Access-1(config)# lldp
```

```
T11-Access-1(config-if)# no lldp receive
T11-Access-1(config-if)# no lldp transmit
T11-Access-1(config-if)# lldp receive
T11-Access-1(config-if)# lldp transmit
```

7. Issue the “**show lldp local device**” command. This will show the information the local device shares/advertises with LLDP messages.

```
T11-Access-2# show lldp local-device

Global Data
===========

Chassis-ID             : 88:3a:40:97:a4:40
System Name            : T11-Access-2
System Description     : Aruba JL668A  FL.10.04.0030
Management Address     : 88:3a:30:97:a4:40
Capabilities Available : Bridge, Router
Capabilities Enabled   : Bridge, Router
TTL                    : 120

Port Based Data
===============

Port-ID             : 1/1/4
Port-Desc           : "1/1/4"
Port Mgmt-Address :
```

```
Port VLAN ID       : 1

Port-ID            : 1/1/28
Port-Desc          : "1/1/28"
Port Mgmt-Address :
Port VLAN ID       : 1


T11-Access-2#
```

What is the “System Description”?

____

What are the available capabilities supported by the system?

____

IMPORTANT: AOS-CX systems have IP routing service enabled by default and cannot be disabled. This means they will automatically populate entries in the Routing Table for whatever IP segment they are configured with in Layer 3 ports (ether physical or logical) and start moving packets at Layer 3 between those segments. IP routing cannot be disabled in these systems.

8. Write down System Name and Chassis ID in **figure 4.2-9**.

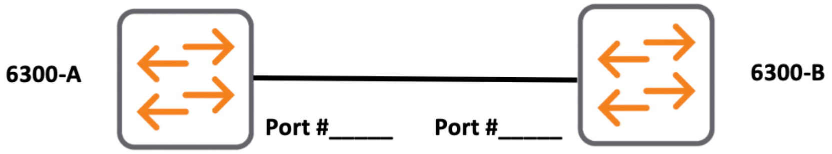

System Name:______________________ **System Name:**______________________

Chassis ID: ________________________ **Chassis ID:** ________________________

Figure 4.2-9: LLDP Discovery.

TIP: You can find a larger copy of this diagram in Appendix 3.

What interfaces are currently running the protocol?

Steps

Access-1

9. Move to **Access-1**.
10. Issue the "**show lldp neighbor-info**" command. You should see only one entry in the output.

```
T11-Access-1# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 1
Total Neighbor Entries Deleted  : 0
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 0

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/28      88:3a:30:97:a4:40  1/1/28        1/1/28        120      T11-Access-2
T11-Access-1#
```

Does the entry match the Chassis-ID and System Name seen in step 8?

What is the local port?

What is the remote port?

11. Try the same command but specify the local interface number at the end of the command.

```
T11-Access-1# show lldp neighbor-info 1/1/28

Port                           : 1/1/28
Neighbor Entries               : 1
Neighbor Entries Deleted       : 0
Neighbor Entries Dropped       : 0
Neighbor Entries Aged-Out      : 0
Neighbor Chassis-Name          : T11-Access-2
Neighbor Chassis-Description   : Aruba JL668A  FL.10.04.0030
Neighbor Chassis-ID            : 88:3a:30:97:a4:40
Neighbor Management-Address    : 88:3a:30:97:a4:40
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled   : Bridge, Router
Neighbor Port-ID               : 1/1/28
Neighbor Port-Desc             : 1/1/28
Neighbor Port VLAN ID          : 1
TTL                            : 120

Neighbor Mac-Phy details
Neighbor Auto-neg Supported    : true
Neighbor Auto-Neg Enabled      : false
Neighbor Auto-Neg Advertised   : Other
Neighbor MAU type              : 10 GIGBASEER

T11-Access-1#
```

NOTE: This version of the command displays the detailed data of the neighbor just like the command, “show lldp local-device” used earlier on Access-2.

12. Finally, run “**show lldp local-device**” on Access-1. Then use the output of this step and the previous step to complete the remaining fields of **figure 4.2-9**.

```
T11-Access-1# show lldp local-device

Global Data
===========

Chassis-ID                : 88:3a:30:98:30:00
System Name               : T11-Access-1
System Description        : Aruba JL668A  FL.10.04.0030
Management Address        : 88:3a:30:98:30:00
```

```
Capabilities Available : Bridge, Router
Capabilities Enabled   : Bridge, Router
TTL                    : 120

Port Based Data
===============

Port-ID           : 1/1/1
Port-Desc         : "1/1/1"
Port Mgmt-Address :
Port VLAN ID      : 1111

Port-ID           : 1/1/3
Port-Desc         : "1/1/3"
Port Mgmt-Address :
Port VLAN ID      : 1111

Port-ID           : 1/1/28
Port-Desc         : "1/1/28"
Port Mgmt-Address :
Port VLAN ID      : 1
```

> **NOTE:** Understanding LLDP and the information it provides can help you verify and troubleshoot Layer 1 communication between devices.

Now that you are sure about which ports are used, you are ready to set the interface descriptions.

13. Set descriptions on both switches' interface 1/1/28.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/28
T11-Access-1(config-if)# description TO_TX-ACCESS-2_PORT-28
T11-Access-1(config-if)# end
```

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface 1/1/28
T11-Access-2(config-if)# description TO_TX-ACCESS-1_PORT-28
T11-Access-2(config-if)# end
```

PC-4

14. Move back to PC-4 and **ping PC-3**'s IP address (**10.X.11.103**).

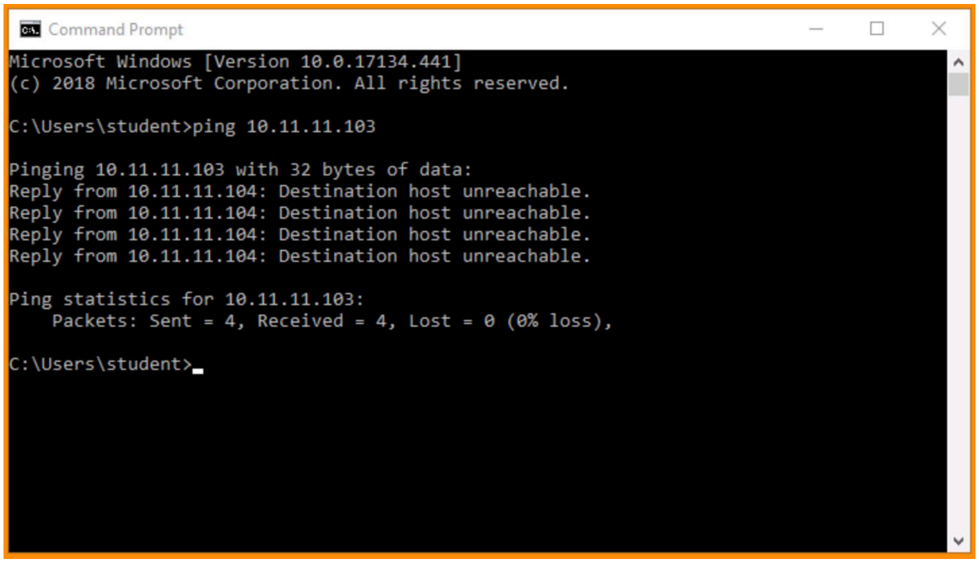

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\student>_
```

Figure 4.2-11: Destination host unreachable.

Was ping successful?

____

Why?

____

ANSWER: Even though a link between both switches has been enabled, ping still fails. In order to better understand why, you should explore the mac-address-table of either switch. Let’s do it on Access-1.

15. Open console session to Access-1 and use the “**show mac-address-table**” command.

```
T11-Access-1# show mac-address-table
MAC age-time             : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                     Port
--------------------------------------------------------------
00:50:56:b1:a9:86    1        dynamic                  1/1/28
```

```
00:50:56:b1:ae:e8    1111      dynamic                  1/1/3
T11-Access-1#
```

TIP: This output may give you more entries than the ones in example above (e.g. PC-1), ignore all but the interfaces to PC-3 and PC-4's.

What Port and VLAN is PC-3 seen on?

____

What Port and VLAN is PC-4 seen on?

____

ANSWER: As you can see both PCs are on different ports (which is expected) and also on different VLANs. PC-4 is seen on VLAN 1 because that is the only VLAN that exists on Access-2, and the only VLAN it forwards in its 1/1/28 interface.

NOTE: As seen in this step, understanding the fundamentals of layer 2 forwarding and exploring the MAC Address table of switches are key tools for troubleshooting the lack of connectivity between two endpoints.

Task 3: Extend Connectivity for VLAN X11

Objectives

After finding the root cause that prevents communication between two endpoints it is time to apply a configuration that solves the issue. You will proceed now to extend VLAN X11 to Access-2 switch.

Steps

Access-1

1. Configure Access-1's interface **1/1/28** as trunk link that permits **VLANs 1 and X11**.

> **NOTE:** Replace the highlighted "X" in the "vlan trunk allowed" command for your student table number, e.g. "vlan trunk allowed 1,411" for table 4, or "vlan trunk allowed 1,1111" for table 11 like in the example below.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/28
T11-Access-1(config-if)# vlan trunk allowed 1,X11
T11-Access-1(config-if)# end
T11-Access-1#
```

2. Display trunk interfaces.

```
T11-Access-1# show interface trunk

----------------------------------------------------------------------
Port    Native VLAN Trunk VLANs
----------------------------------------------------------------------
1/1/28  1           1,1111
T11-Access-1#
```

Access-2

3. Move to **Access-2.**
4. Create **VLAN X11** and name it **EMPLOYEES**.

> **NOTE:** Replace the highlighted “X” in the vlan command for your student table number, e.g. “vlan 411” for table 4, or “vlan 1111” for table 11 like in the example below.

```
T11-Access-2# configure terminal
T11-Access-2(config)# vlan X11
T11-Access-2(config-vlan-1111)# name EMPLOYEES
T11-Access-2(config-vlan-1111)# exit
```

5. Configure **Access-2’s** interface 1/1/28 as trunk link that permits **VLANs 1 and X11**.

```
T11-Access-2(config)# interface 1/1/28
T11-Access-2(config-if)# vlan trunk allowed 1,X11
T11-Access-2(config-if)# exit
```

6. Last configure interface **1/1/4** as access port in **VLAN X11**.

```
T11-Access-2(config)# interface 1/1/4
T11-Access-2(config-if)# vlan access X11
T11-Access-2(config-if)# end
```

7. Confirm **VLAN X11** is now member of ports **1/1/1 and 1/1/28**.

```
T11-Access-2# show vlan X11

--------------------------------------------------------------------------------
VLAN  Name                                Status   Reason                 Type
Interfaces
--------------------------------------------------------------------------------
1111  employees                           up       ok                     static
1/1/4,1/1/28
T11-Access-2#
```

8. Display trunk interfaces. You should have only one trunk port.

```
T11-Access-2# show interface trunk

----------------------------------------------------------------------
Port    Native VLAN Trunk VLANs
----------------------------------------------------------------------
1/1/28  1           1,1111
T11-Access-2#
```

9. Move back to PC-4 and ping **PC-3**’s IP address (**10.X.11.103**).

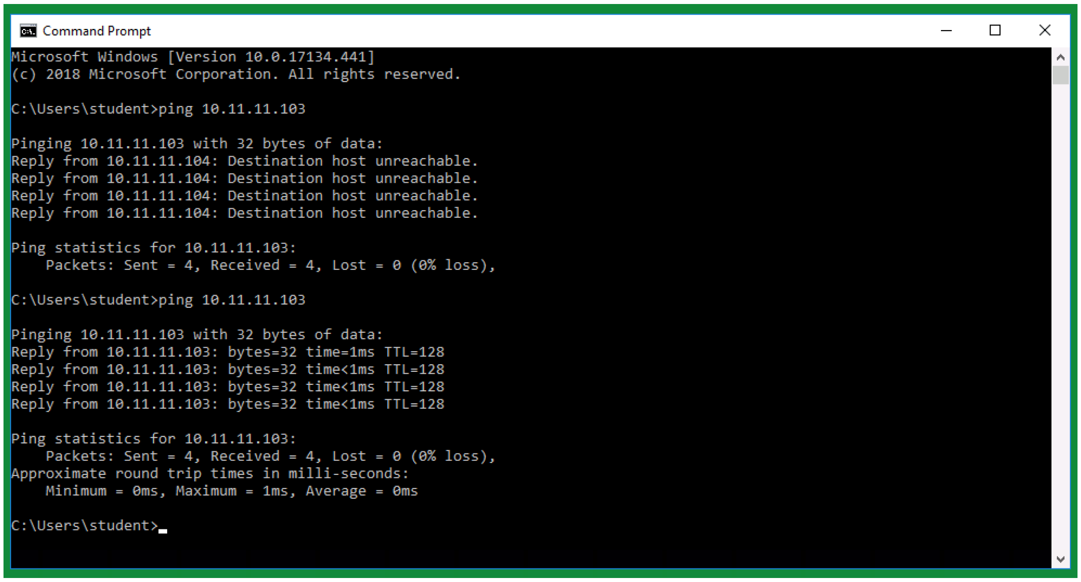

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.
Reply from 10.11.11.104: Destination host unreachable.

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.103: bytes=32 time=1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128
Reply from 10.11.11.103: bytes=32 time<1ms TTL=128

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\student>_
```

Figure 4.2-12: Ping successful.

Was ping successful?

____

Let’s now explore the MAC address tables of both switches and trace the MAC addresses of each station in order to confirm they are learned in the expected ports and VLANs.

Access-1 and Access-2

10. Display the mac address table of both Access-1 and Access-2.

```
T11-Access-2# show mac-address-table
MAC age-time                : 300 seconds
```

```
Number of MAC addresses : 4

MAC Address          VLAN     Type                      Port
--------------------------------------------------------------
00:50:56:b1:a9:86    1111     dynamic                   1/1/4
00:50:56:b1:ae:e8    1111     dynamic                   1/1/28
T11-Access-2#
```

```
T11-Access-1# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                      Port
--------------------------------------------------------------
00:50:56:b1:a9:86    1111     dynamic                   1/1/28
00:50:56:b1:ae:e8    1111     dynamic                   1/1/3
T11-Access-1#
```

11. With the information shown please fill out the fields on **figure 4.2-13**.

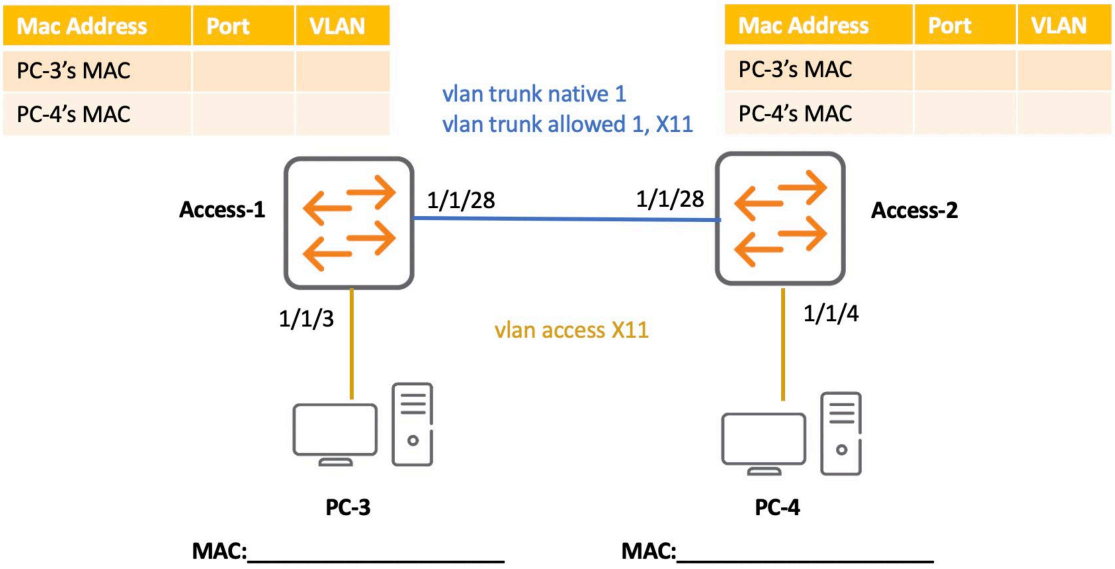

Mac Address	Port	VLAN
PC-3's MAC		
PC-4's MAC		

Mac Address	Port	VLAN
PC-3's MAC		
PC-4's MAC		

Figure 4.2-13: MAC address tables.

TIP: You can find a larger copy of this diagram in Appendix 3.

Task 4: Save Your Configurations

Objectives

You will now save your configurations and create checkpoints. Remember, final lab checkpoints may be used in later activities.

Steps

Access-1 and Access-2

1. Save the current Access switches' configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab4-2_final**.

```
T11-Access-1# copy running-config checkpoint Lab4-2_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab4-2_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 4.2!

AOS-CX Switching Fundamentals

Lab 4.3: Add a Core Switch to the Topology

Overview

After a few months in business, BigStartup seems to have a promising forecast. Sales are growing and more employees are being hired. The company is urgently investigating renting the West Wing of the floor. Management is considering the implications of expansion and what effect it will have on the network.

They have approached you for advice and you have recommended the insertion of a Core switch, following a 2-tier design that can assure future growth with no complexity (instead of a daisy chain-based topology). You suggest an 8325 AOS-CX switch, which assures a consistent OS across the board, high port density, unprecedent throughput and no blocking switching. While management agrees with your recommendation and can budget for the new gear, it turns out that the building owner, Cheap4Rent, also offers some degree of network services for all their tenants.

Cheap4Rent offers to include the same 8325 AOS-CX switch in the lease. This permits the company to save capital and invest in other assets such as servers, IP telephony, video surveillance, etc.

BigStartup is the first tenant to be offered the Core Switching service and to facilitate the integration, they are giving you limited network operations access over SSH and will allow you to use the default VRF for now.

Objectives

After completing this lab, you will be able to:

- Deploy a Core Switch to the topology
- Configure uplinks as trunk ports by enabling 802.1Q
- Add a new VLAN for another users' type
- Enable DHCP server on Access-1

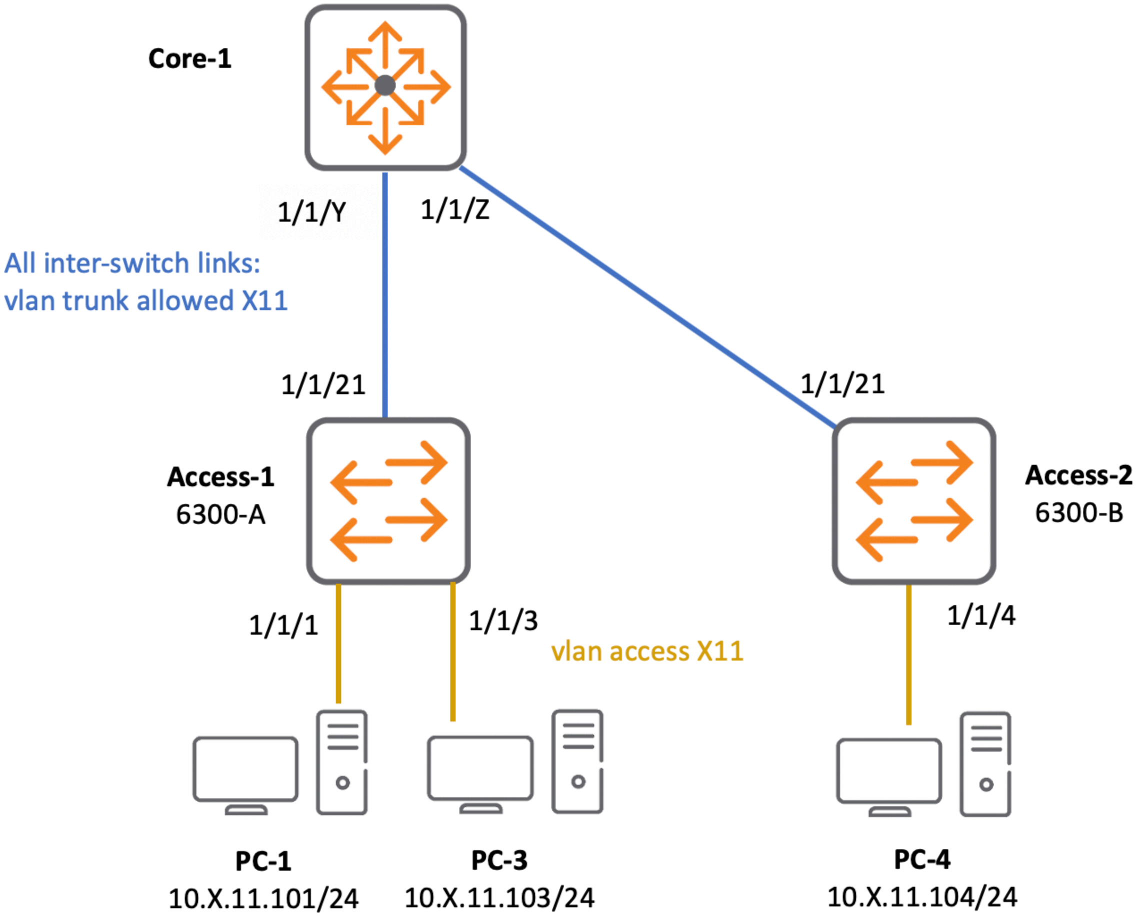

Figure 4.3-1: Lab Topology

Task 1: Add Core-1 to the Topology.

Objectives

In this task, you will change the switching topology and enable ports on the Access switches that have been connected to the 8325 AOS-CX Core Switch that resides in the Building's MDF. You will also configure the core switch side of the links and validate the topology.

Even though 8300 platforms come with disabled routed ports by default, Cheap4Rent has turned the Core ports on and made them switch interfaces. They have provided ethernet wire drops for establishing Layer 1 connectivity between the Core and Access switches.

Steps

Access-1 and Access-2

1. Disable the link between **Access-1** and **Access-2**.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/28
T11-Access-1(config-if)# shutdown
```

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface 1/1/28
T11-Access-2(config-if)# shutdown
```

Access-1

2. Move back to Access-1.
3. Allow **VLAN X11** as a tagged member of port **1/1/21** and enable the interface.

```
T11-Access-1(config)# interface 1/1/21
T11-Access-1(config-if)# vlan trunk allowed X11
T11-Access-1(config-if)# no shutdown
```

TIP: You were told by the Cheap4Rent team that your switches were connected on ports 1/1/3 and 1/1/6 on the Core side, nonetheless you know from experience that it is always better to verify third-party technical information using LLDP.

4. Use the "**show lldp neighbor-info"** command, to validate the port Access-1 is connected to.

```
T11-Access-1(config-if)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 2
Total Neighbor Entries Deleted  : 1
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 1

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/21      90:20:c2:bc:ed:00  1/1/16        1/1/16        120      Core-1
T11-Access-1(config)#
```

TIP: This output may still show Access-2 on port 1/1/28. That would be an old entry that is about to age out.

Was the information given by Rent4Cheap accurate?

____

5. Use the information in the output for drawing the connection and setting the right ports numbers on **figure 4.3-2**. Also set the proper description on port 1/1/21.

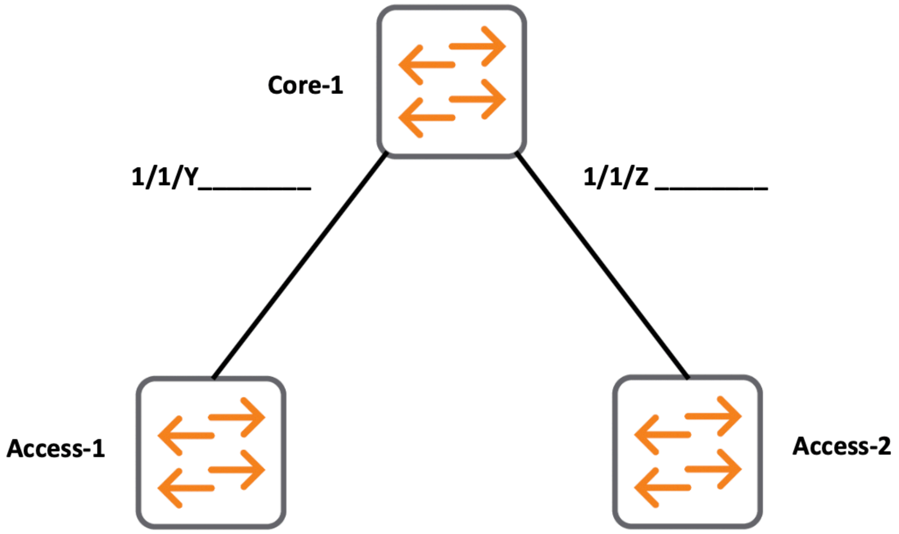

Figure 4.3-2: LLDP Discovery

TIP: You can find a larger copy of this diagram in Appendix 3.

NOTE: In the subsequent steps, downlinks from Core switches to Access-1 will be referred as ports “**Y**” and downlinks to Access-2 will be referred as ports “**Z**”. Ask your instructor if you have any questions.

```
T11-Access-1(config-if)# description TO_CORE-1_PORT-Y
T11-Access-1(config-if)#
```

6. Move back to **Access-2** and **repeat steps 3 to 5**. Do not forget to draw the connections in **figure 4.3-2**.

Access-2

```
T11-Access-2(config)# interface 1/1/21
T11-Access-2(config-if)# vlan trunk allowed X11
T11-Access-2(config-if)# no shutdown
```

```
T11-Access-2(config-if)# show lldp neighbor-info

LLDP Neighbor Information
=========================
Total Neighbor Entries          : 2
Total Neighbor Entries Deleted  : 1
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 1

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/21      90:20:c2:bc:ed:00  1/1/37        1/1/37        120      Core-1
T11-Access-2(config-if)#
```

```
T11-Access-2(config-if)# description TO_CORE-1_PORT-Z
T11-Access-2(config-if)# end
```

Just as a sanity check you will connect to Core-1 and confirm the connection status on that device. To access it you will connect to PC-1 and use it as a “jump host” running an SSH session to Core-1’s IP address.

TIP: PC-1 has two Lab related ethernet connections, “LAB NIC” and “OOBM” (Out of Band Management). You will access Core-1 using the second one as shown in the figure.

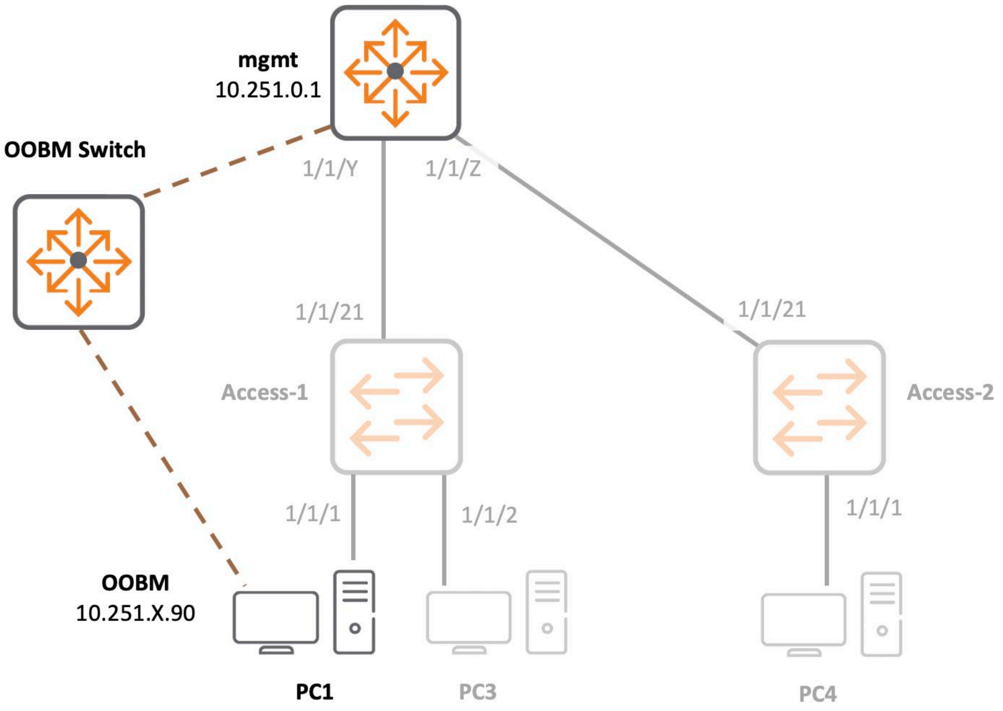

Figure 4.3-3: Using OOBM network.

PC-1

7. Access the **PC-1**.
8. Open Putty. You will find saved sessions to Core-1 and other three devices.

TIP: Putty should have Saved Sessions to Core-1 and Core-2, you could use those as a shortcut.

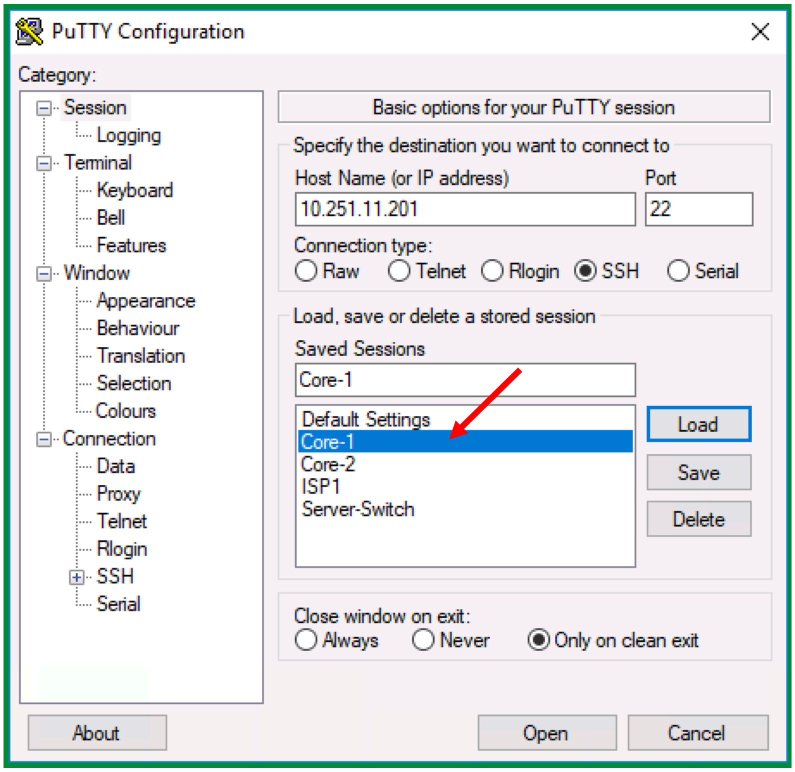

Host Name (or IP address): 10.251.11.201
Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Saved Sessions: Core-1

Default Settings
Core-1
Core-2
ISP1
Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 4.3-4: Putty

9. Double click **Core-1** saved session.

Core-1 (via PC-1)

10. Login using **cxfX/aruba123**

NOTE: Replace the highlighted "X" with your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

11. Define the height of the page as 40 lines.

```
Core-1# page 40
Core-1#
```

12. Type "**show lldp neighbor-info | include TX**"

NOTE: The information will be filtered out, listing only the lines that include the "T**X**" string.

NOTE: Replace the highlighted "X" for your student table number, e.g. **info | include T4** for table 4, or **| include T11** for table 11 like in the example below.

```
Core-1# show lldp neighbor-info | include TX
1/1/16      88:3a:30:98:30:00  1/1/21        1/1/21        120      T11-Access-1
1/1/37      88:3a:30:97:a4:40  1/1/22        1/1/22        120      T11-Access-2
Core-1#
```

NOTICE: The pipe (|) command filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.

Strings of characters that follow the filtering tool (e.g. "T4" or "T11" in example above) are case sensitive. Typing the wrong capitalization may lead to the absence of output.

Does the output match what you recorded on figure 4.3-2?

13. Create **VLAN X11** and name it **TX_EMPLOYEES.**

NOTE: Replace the highlighted “X” for your student table number, e.g. **vlan 411 / name T4_EMPLOYEES** for table 4, or **vlan 1111 / name T11_EMPLOYEES** for table 11 like in the example below.

```
Core-1# configure terminal
Core-1(config)# vlan X11
Core-1(config-vlan-1111)# name TX_EMPLOYEES
Core-1(config-vlan-1111)# exit
Core-1(config)#
```

NOTICE: Command based authorization is enabled on all SSH sessions you will run in this training lab. This means that every command you type on SSH sessions will be validated with a list of permitted commands. If the command you type is not in the list, you will get an error messages like the following:

```
Core-2(config)# vlan 1999
Cannot execute command. Command not allowed.
Core-2(config)#
```

14. Access port **1/1/Y**, then set the **TO_TX-ACCESS-1_PORT-21** description and make the interface switch port and a trunk interface member of **VLAN X11**.

NOTE: Replace the highlighted “Y” for the port number of the downlink that connects to Access-1 as recorded on figure 4.3-2. Also replace the highlighted “X”s for your student table number.

```
Core-1(config)# interface 1/1/Y
Core-1(config-if)# description TO_TX-ACCESS-1_PORT-21
Core-1(config-if)# vlan trunk allowed X11
Core-1(config-if)#
```

15. Move to port **1/1/Z**, then repeat **step 11** using **TO_TX-ACCESS-2_PORT-21** as description

NOTE: Replace the highlighted “Z” for the port number of the downlink that connects to Access-1 as recorded on figure 4.3-2. Also replace the highlighted “X” for your student table number.

```
Core-1(config)# interface 1/1/Z
Core-1(config-if)# description TO_TX-ACCESS-2_PORT-21
Core-1(config-if)# vlan trunk allowed X11
Core-1(config-if)# end
```

PC-1

16. From PC-1 **ping PC-4** (10.X.11.104). Ping should be successful.

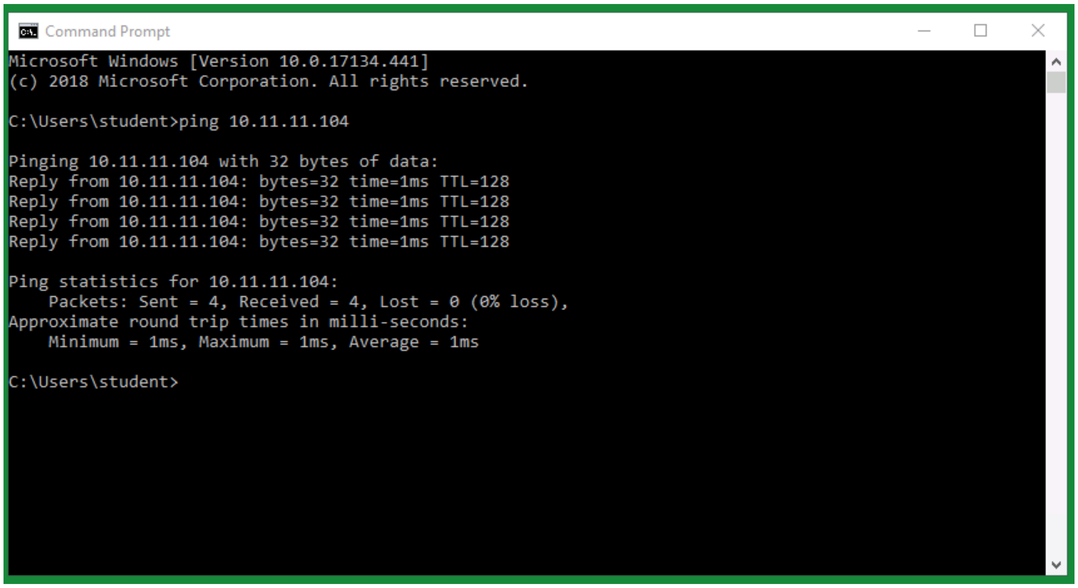

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.11.104

Pinging 10.11.11.104 with 32 bytes of data:
Reply from 10.11.11.104: bytes=32 time=1ms TTL=128
Reply from 10.11.11.104: bytes=32 time=1ms TTL=128
Reply from 10.11.11.104: bytes=32 time=1ms TTL=128
Reply from 10.11.11.104: bytes=32 time=1ms TTL=128

Ping statistics for 10.11.11.104:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 1ms, Average = 1ms

C:\Users\student>
```

Figure 4.3-5: Ping successful

Core-1 (via PC-1)

17. **OPTIONAL** - You can display the MAC address table to see what ports Core-1 learned the clients' MAC addresses from, which are the ports it uses for forwarding traffic to them at Layer 2.

```
Core-1# show mac-address-table vlan X11
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                          Port
--------------------------------------------------------------
00:50:56:b1:ae:e8    1111     dynamic                       1/1/16
00:50:56:b1:37:67    1111     dynamic                       1/1/16
00:50:56:b1:a9:86    1111     dynamic                       1/1/37
Core-1#
```

Task 2: Adding a Second VLAN.

Objectives

After more hiring, BigStartup is now interested in improving privacy and traffic separation between regular employees and managers. They are asking you if there is any way you can achieve that with networking devices they already have. You can improve privacy and traffic separation by adding another VLAN.

The next steps will be focused on creating VLAN X12 for managers across all switches and moving PC-1 and PC-4 into that broadcast domain.

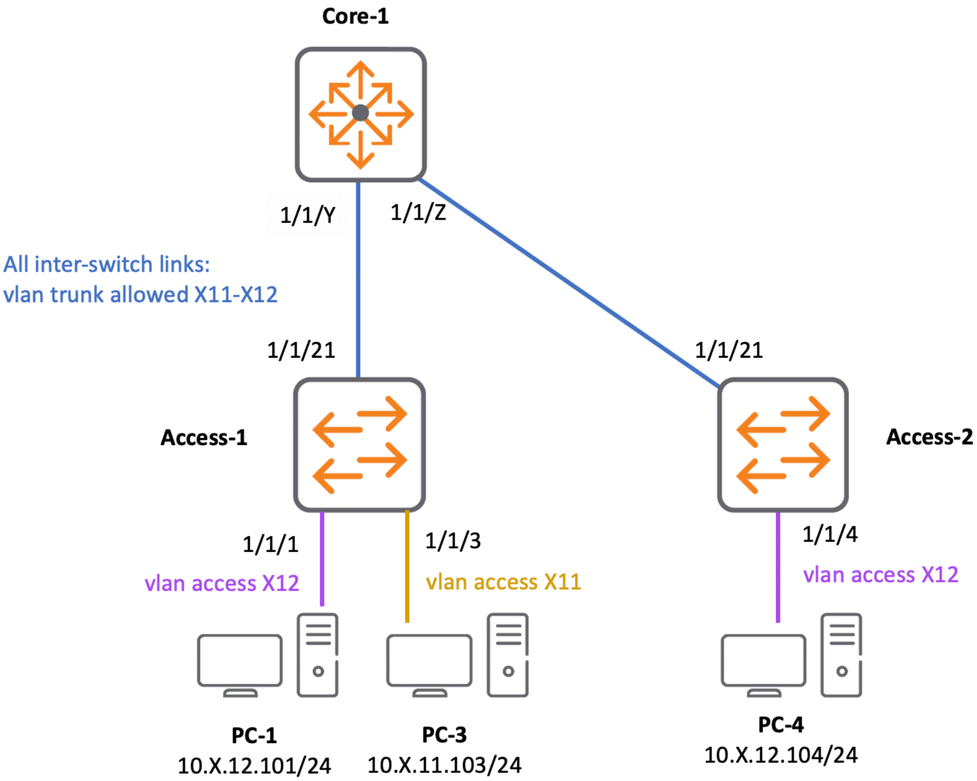

Figure 4.3-6: Lab4.3-Task 2 Logical Topology.

Steps

Access-1

1. Open a console connection to the Access-1. Login using **admin** and **no password**.
2. Create **VLAN X12** and name it **MANAGERS**, then apply it on port **1/1/21**.

NOTE: Replace the highlighted "X" for your student table number.

```
T11-Access-1# configure terminal
T11-Access-1(config)# vlan X12
T11-Access-1(config-vlan-1112)# name MANAGERS
T11-Access-1(config-vlan-1112)# exit
T11-Access-1(config)#
T11-Access-1(config)# interface 1/1/21
T11-Access-1(config-if)# vlan trunk allowed X12
T11-Access-1(config-if)# exit
T11-Access-1(config)#
```

3. Use the "**show vlan**" command to see the new added VLAN and the port members.

```
T11-Access-1(config)# show vlan

---------------------------------------------------------------------------------
VLAN  Name                    Status  Reason                Type      Interfaces
---------------------------------------------------------------------------------
1     DEFAULT_VLAN_1.         down    no_member_forwarding  default
1/1/2,1/1/4-1/1/20,1/1/22-1/1/28
1111  EMPLOYEES               up      ok                    static
1/1/1,1/1/3,1/1/21,1/1/28
1112  MANAGERS                up      ok                    static    1/1/21
T11-Access-1(config)#
```

Access-2

4. Open a console connection to the **Access-2.** Login using **admin** and **no password**.
5. Repeat **step 2**.

```
T11-Access-2# configure terminal
T11-Access-2(config)# vlan X12
T11-Access-2(config-vlan-1112)# name MANAGERS
T11-Access-2(config-vlan-1112)# exit
```

```
T11-Access-2(config)#
T11-Access-2(config)# interface 1/1/21
T11-Access-2(config-if)# vlan trunk allowed X12
T11-Access-2(config-if)# exit
T11-Access-2(config)#
```

Core-1 (via PC-1)

6. Move back to **Core-1** SSH session.
7. Create **VLAN X12**, name it **TX_MANAGERS.**

> **NOTE:** Replace the highlighted “X” for your student table number, e.g. **vlan 411 / name T4_MANAGERS** for table 4, or **vlan 1111 / name T11_MANAGERS** for table 11 like in the example below.

```
Core-1# configure terminal
Core-1(config)# vlan X12
Core-1(config-vlan-1112)# name TX_MANAGERS
Core-1(config-vlan-1112)# exit
```

8. Apply **VLAN X12** on **port 1/1/Y**.

> **NOTE:** Replace the highlighted “X” for your student table number, also replace highlighted “Y” for the port number of the downlink that connects to Access-1 as recorded on figure 4.3-2.

```
Core-1(config)# interface 1/1/Y
Core-1(config-if)# vlan trunk allowed X12
Core-1(config-if)#
```

9. Repeat **step 8** on interface **1/1/Z**.

> **NOTE:** Replace highlighted “Z” for the port number of the downlink that connects to Access-2 as recorded on figure 4.3-2.

All switches have VLANs X11 and X12 now, and they have been assigned in all switch to switch links. Now you will move PC1 and PC4 into VLAN X12 and test connectivity.

Access-1

10. Move to **Access-1.**
11. Make interface **1/1/1** an access port on **VLAN X12.**

> **NOTE:** Replace the highlighted "X" for your student table number.

```
T11-Access-1(config)# interface 1/1/1
T11-Access-1(config-if)# vlan access X12
T11-Access-1(config-if)#
```

Access-2

12. Move to **Access-2.**
13. Make interface **1/1/4** an access port on **VLAN X12.**

```
T11-Access-2(config)# interface 1/1/4
T11-Access-2(config-if)# vlan access X12
T11-Access-2(config-if)#
```

You will now change the IP segment where PC-1 and PC-4 belong.

PC-1

14. Access PC-1 and change the "**Lab NIC**" IP address to **10.X.12.101/24**

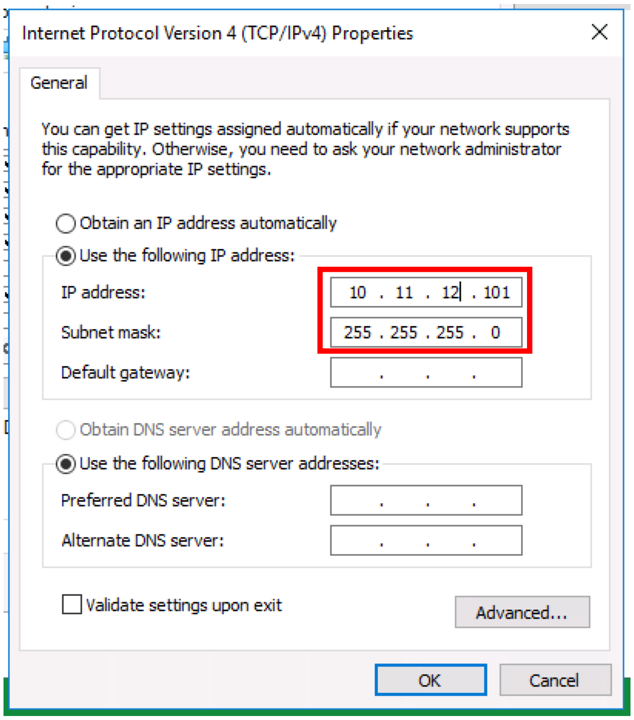

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 12 . 101

Subnet mask: 255 . 255 . 255 . 0

Default gateway: . . .

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 4.3-7: PC-1's IP address setting.

15. Use the "**ipconfig -all**" command and confirm the client is using the new IP address.

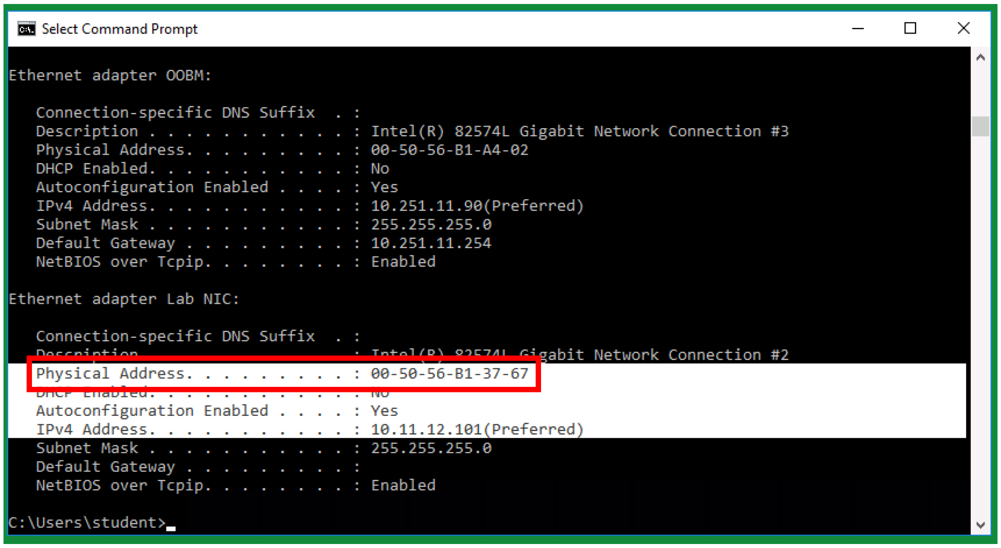

Select Command Prompt

```
Ethernet adapter OOBM:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection #3
   Physical Address. . . . . . . . . : 00-50-56-B1-A4-02
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 10.251.11.90(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.251.11.254
   NetBIOS over Tcpip. . . . . . . . : Enabled

Ethernet adapter Lab NIC:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection #2
   Physical Address. . . . . . . . . : 00-50-56-B1-37-67
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 10.11.12.101(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   NetBIOS over Tcpip. . . . . . . . : Enabled

C:\Users\student>_
```

Figure 4.3-8: PC-1 network settings

What is the NIC MAC address?

PC-4

16. Access PC-4 and change the “**Lab NIC**” IP address to **10.X.12.104/24**

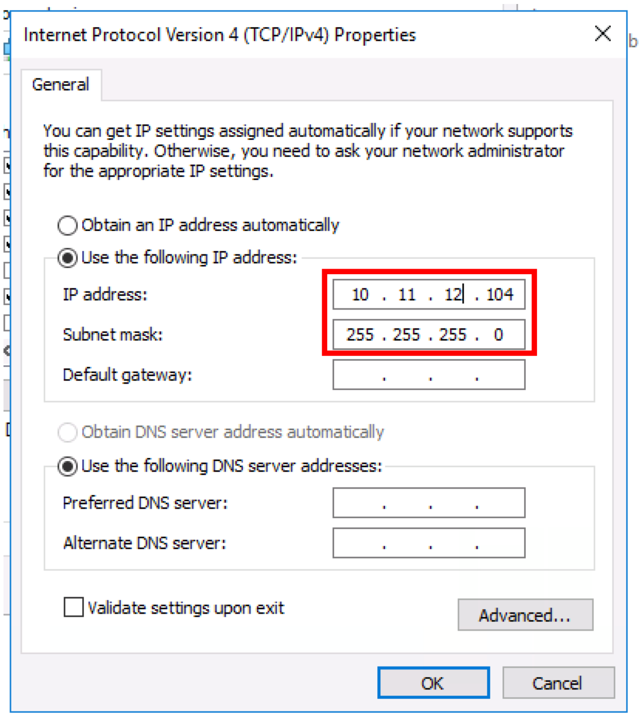

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 12 . 104

Subnet mask: 255 . 255 . 255 . 0

Default gateway: . . .

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 4.3-9: PC-4’s IP address setting.

17. Ping **PC-1** (**10.X.12.101**).

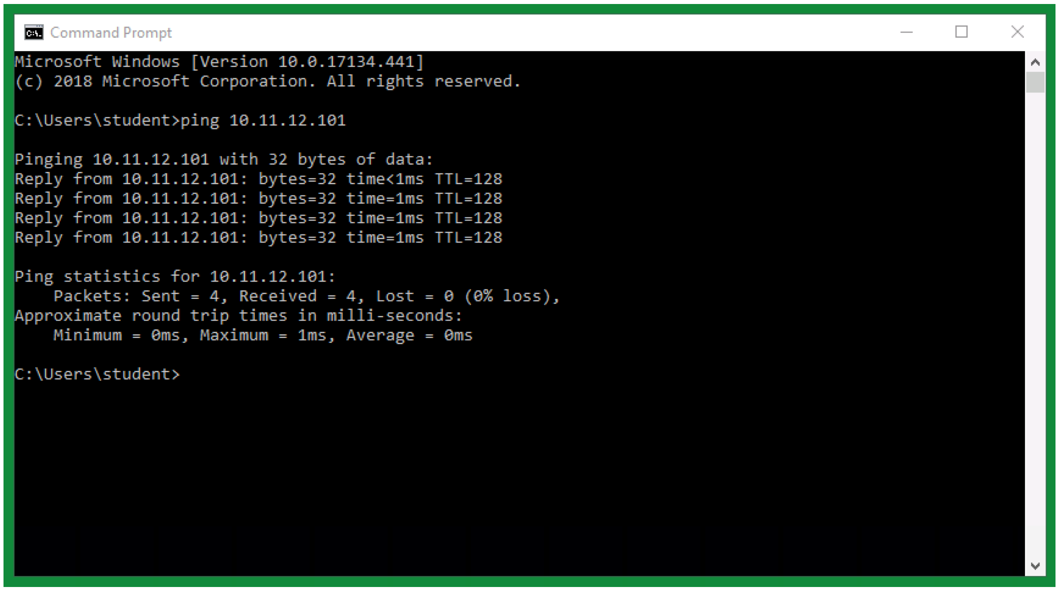

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.12.101

Pinging 10.11.12.101 with 32 bytes of data:
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128

Ping statistics for 10.11.12.101:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\student>
```

Figure 4.3-10: Ping successful.

Was ping successful?

____

18. Ping **PC-3** (**10.X.11.103**).

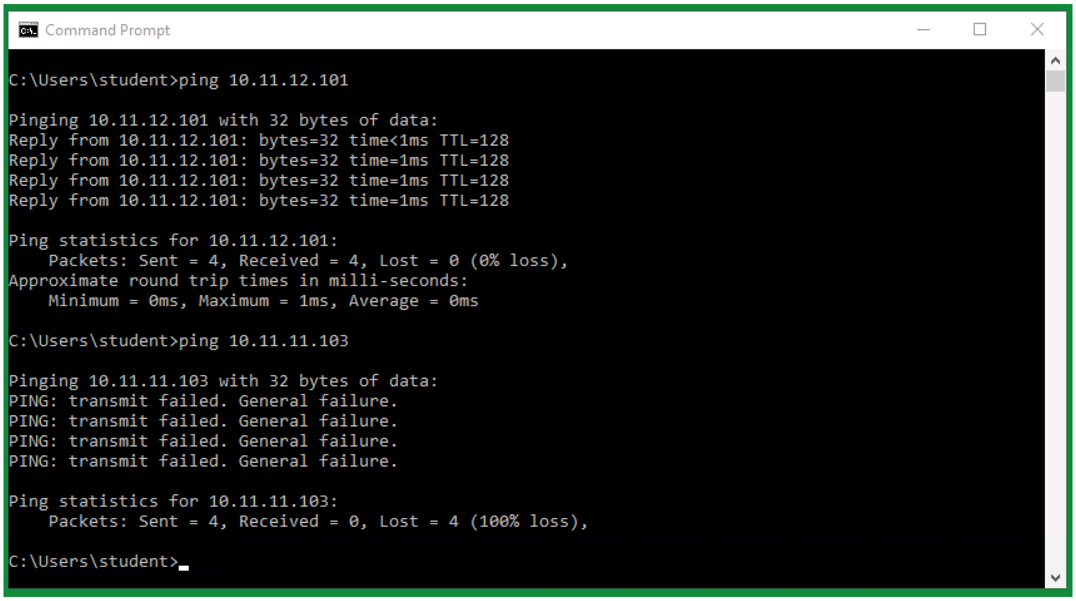

```
C:\Users\student>ping 10.11.12.101

Pinging 10.11.12.101 with 32 bytes of data:
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128

Ping statistics for 10.11.12.101:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
PING: transmit failed. General failure.
PING: transmit failed. General failure.
PING: transmit failed. General failure.
PING: transmit failed. General failure.

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>_
```

Figure 4.3-11: Ping unsuccessful.

Was ping successful?

____

ANSWER: Pinging PC-3 will fail because it is now in a different network.

19. Display the ARP Table using the "**arp -a**" command and look for the **10.X.12.101** entry.

TIP: You can use the filtered version of this command "arp -a -N 10.X.12.104" for only displaying entries associated with "Lab NIC" interface.

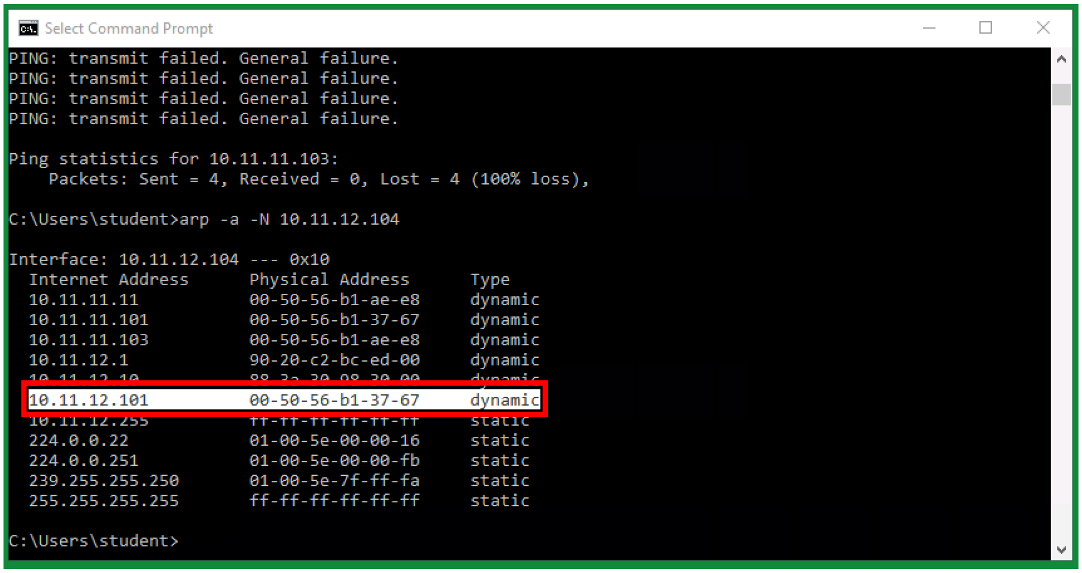

```
PING: transmit failed. General failure.
PING: transmit failed. General failure.
PING: transmit failed. General failure.
PING: transmit failed. General failure.

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>arp -a -N 10.11.12.104

Interface: 10.11.12.104 --- 0x10
  Internet Address      Physical Address      Type
  10.11.11.11           00-50-56-b1-ae-e8     dynamic
  10.11.11.101          00-50-56-b1-37-67     dynamic
  10.11.11.103          00-50-56-b1-ae-e8     dynamic
  10.11.12.1            90-20-c2-bc-ed-00     dynamic
  10.11.12.10           88-3a-30-98-30-00     dynamic
  10.11.12.101          00-50-56-b1-37-67     dynamic
  10.11.12.255          ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  239.255.255.250       01-00-5e-7f-ff-fa     static
  255.255.255.255       ff-ff-ff-ff-ff-ff     static

C:\Users\student>
```

Figure 4.3-12: PC-4's ARP table.

Is the MAC address in the entry the same you recorded in **step 15**?

____

NOTE: You might also see a 10.X.11.101 entry associated with the same MAC. That is an old record from the time PC-1 and PC-4 were both in VLAN X11, this entry will eventually expire.

Access-1

20. Move to **Access-1**.
21. Display the MAC address table. You will see one entry associated with **VLAN X11** and another with **VLAN X12**.

```
T11-Access-1# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                      Port
--------------------------------------------------------------
00:50:56:b1:ae:e8    1111     dynamic                   1/1/3
00:50:56:b1:37:67    1112     dynamic                   1/1/1
00:50:56:b1:a9:86    1112     dynamic                   1/1/21
T11-Access-1#
```

NOTE: If you do not get an entry mapped to port 1/1/3, artificially generate some traffic on PC-3 to let Access-1 re-learn its MAC address again. A single ping to 10.X.11.101 is enough. It will work even if the ping is unsuccessful.

Task 3: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

Access-1, Access-2 and Core-1 (via PC-1)

1. Save the current Access switches and Core-1 configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

Access-1 and Access-2

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab4-3_final**.

```
T11-Access-1# copy running-config checkpoint Lab4-3_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab4-3_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 4.3!

AOS-CX Switching Fundamentals

Lab 5.1: Rapid Spanning Tree Protocol

Overview

Your Core switch integration has proven successful and the network is more scalable, however, experience tells you that a single Core switch is a single point of failure. If an uplink or the Core itself goes down all business operations will be disrupted. During a conversation you share this concern with BigStartup management. A formal request for a second 8325 switch was sent to Rent4Cheap Properties who agreed to supply the second unit and modify the lease. A few weeks later the switch arrived and was connected to Core-1.

BigStartup has notified you the additional Core switch is operational and has asked you to complete the integration.

Objectives

After completing this lab, you will be able to:

- Add a redundant core switch
- Enable redundant links
- Verify the spanning tree functionality
- Find the Root bridge
- Discover the CST topology

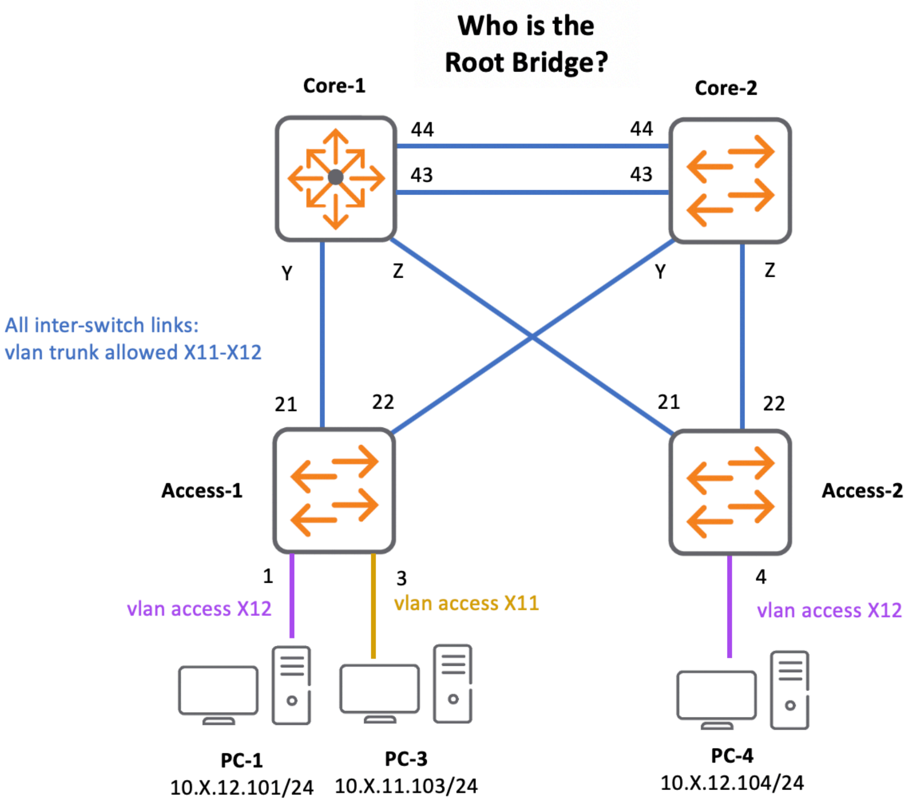

Figure 5.1-1: Lab Topology

Task 1: Add the Redundant Core Switch and Redundant Links

Objectives

In this task you will add a fourth component to the topology: Core-2. First you will make sure that the Core and Access switches are running Spanning Tree. Next, you will prepare port 1/1/22 on both Access switches to act as uplinks to Core-2 and enable them.

Finally, you will confirm that connectivity between hosts is still in place.

Steps

PC-1

1. Access PC-1.
2. Open Putty and open a SSH session to Core-1 (**10.251.0.1**) and login with **cxfX/aruba123**

NOTE: Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

TIP: Putty should have Saved Sessions to Core-1 and Core-2, you can use these as shortcuts.

Core-1 (via PC-1)

3. Define the height of the page to 40 lines.

```
Core-1# page 40
Core-1#
```

4. Confirm STP is active.

```
Core-1# configure terminal
Core-1(config)# show spanning-tree
Spanning tree status       : Enabled Protocol: MSTP

MST0
  Root ID    Priority   : 0
             MAC-Address: 90:20:c2:bc:ed:00
             This bridge is the root
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority  : 0
             MAC-Address: 90:20:c2:bc:ed:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Disabled       Blocking     20000          128        point_to_point
1/1/2        Disabled       Blocking     20000          128        point_to_point
                            ←---- output omitted ---
```

PC-1

5. Open Putty and open a SSH session to Core-2 (10.251.0.2), and login with: **cxfX/aruba123**

Core-2 (via PC-1)

6. Confirm STP is active.

```
Core-2(config)# show spanning-tree | include Spanning
Spanning tree status       : Enabled Protocol: MSTP
Core-2(config)#
```

Access-1 and Access-2

7. Repeat **step 6** on **Access-1 and Access-2**.

NOTE: The information will be filtered out, listing only the lines that include the "Spanning" string.

```
T11-Access-1# show spanning-tree | include Spanning
Spanning tree status        : Enabled Protocol: MSTP
T11-Access-1#
```

NOTICE: The pipe (|) command filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.

Strings of characters that follow the filtering tool (e.g. "T4" or "T11" in example above) are case sensitive. Typing the wrong capitalization may lead to the absence of output.

```
T11-Access-2# show spanning-tree | include Spanning
Spanning tree status        : Enabled Protocol: MSTP
T11-Access-2#
```

IMPORTANT: Spanning Tree Protocol is enabled by default on 6300s, however in the case of the **8325s** its initial configuration state is **disabled**.

Once enabled, default STP mode is Multiple-Instance Spanning Tree (MST).

IMPORTANT: MST0 relates to instance 0 of MST, this instance is used for interoperating with RSTP switches and MST switches in other regions and to create the Common Spanning Tree (CST): a single Spanning Tree topology for all VLANs.

As a sanity check you will connect to Core-1 and confirm the connections from that device.

Access-1

8. Move back to Access-1.
9. Allow **VLANs X11 and X12** on interface **1/1/22** and enable the port.

NOTE: Replace the highlighted “X” for your student table number.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/22
T11-Access-1(config-if)# vlan trunk allowed X11-X12
T11-Access-1(config-if)# no shutdown
T11-Access-1(config-if)#
```

10. On the Access switch use LLDP to discover which Core-2 remote port is connected to interface **1/1/22**. This will be port **1/1/Y**.

```
T11-Access-1(config-if)# show lldp neighbor-info

LLDP Neighbor Information
=========================

                           ←---- output omitted ---→

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/21      90:20:c2:bc:ed:00  1/1/16        1/1/16        120      Core-1
1/1/22      90:20:c2:bc:3f:00  1/1/16        1/1/16        120      Core-2
```

11. Apply a description to the port.

```
T11-Access-1(config-if)# description TO_CORE-2_PORT-Y
T11-Access-1(config-if)#
```

TIP: Core-2 uses the same port numbers for connecting to your Access Switches as Core-1 did in previous labs (represented by Y and Z).

Access-2

12. Move to Access-2 and repeat **steps 9 to 11**. The remote port that interface **1/1/22** is connected to on the Core-2 side will be **1/1/Z**.

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface 1/1/22
T11-Access-2(config-if)# vlan trunk allowed X11-X12
```

```
T11-Access-2(config-if)# no shutdown
```

```
T11-Access-2(config-if)# show lldp neighbor-info

LLDP Neighbor Information
=========================

                           ←---- output omitted ---→


LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/21      90:20:c2:bc:ed:00  1/1/37        1/1/37        120      Core-1
1/1/22      90:20:c2:bc:3f:00  1/1/37        1/1/37        120      Core-2
```

```
T11-Access-2(config-if)# description TO_CORE-2_PORT-Z
T11-Access-2(config-if)#
```

You have prepared the Access switches uplinks, now you will prepare the connections between the cores and their downlinks.

Core-1

13. Use LLDP to discover the ports used for the connection to Core-2. Use a filtered version of this command to display relevant output only.

```
Core-1# show lldp neighbor-info | exclude 1/21

                           ←---- output omitted ---→

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/43      90:20:c2:bc:3f:00  1/1/43        1/1/43        120      Core-2
1/1/44      90:20:c2:bc:3f:00  1/1/44        1/1/44        120      Core-2
```

What are the Core-2 local ports?

____

What are Core-2 remote ports?

____

14. Move to ports **1/1/43 and 1/1/44** and make each port a trunk interface that allows **VLANs X11 and X12**.

NOTE: Replace the highlighted “X” for your table number.

```
Core-1(config)# interface 1/1/43-1/1/44
Core-1(config-if-<1/1/43-1/1/44>)# vlan trunk allowed X11-X12
Core-1(config-if-<1/1/43-1/1/44>)# exit
Core-1(config)#
```

NOTICE: If when applying configuration above you get the following error message:

```
Core-1(config-if-<1/1/43-1/1/44>)# vlan trunk allowed X11-X12
Operation not allowed on an interface part of a LAG (lag10).
```

This implies that your instructor has already run “Lab - 6.1 - Link demonstration”. This means that interface LAG 10 is replacing ports 43 and 44. You can continue by configuring the lag interface instead of ports 1/1/43 and 1/1/44. Please ask your instructor for more information.

Core-2

15. Open the SSH session of **Core-2**.
16. Define the height of the page to 40 lines.

```
Core-2# page 40
Core-2#
```

17. Create **VLAN X11** and name it **TX-EMPLOYEE.**

NOTE: Replace the highlighted “X” for your student table number, e.g. **vlan 411 / name T4_EMPLOYEES** for table 4, or **vlan 1111 / name T11_MPLOYEES** for table 11 like in the example below.

```
Core-2# configure terminal
Core-2(config)# vlan X11
```

```
Core-2(config-vlan-1111)# name TX_EMPLOYEES
Core-2(config-vlan-1111)# exit
Core-2(config)#
```

18. Create **VLAN X12** and name it **TX_MANAGERS**.

```
Core-2# configure terminal
Core-2(config)# vlan X12
Core-2(config-vlan-1112)# name TX_MANAGERS
Core-2(config-vlan-1112)# exit
Core-2(config)#
```

19. Access port **1/1/Y**. Make the description **TO_TX-ACCESS-1_PORT-22** and make the interface a trunk interface that allows **VLANs X11 and X12**.

> **NOTE:** Replace the highlighted “Y” for the port number of the downlink that connects to Access-1 as discovered in step 3.

```
Core-2(config)# interface 1/1/Y
Core-2(config-if)# description TO_TX-ACCESS-1_PORT-22
Core-2(config-if)# vlan trunk allowed X11-X12
Core-2(config-if)#
```

20. Move to port **1/1/Z**, then set the **TO_TX-ACCESS-2_PORT-22** description and make the interface a trunk interface that allows **VLANs X11 and X12**.

> **NOTE:** Replace the highlighted “Z” for the port number of the downlink that connects to Access-1 as discovered in step 5.

```
Core-2(config)# interface 1/1/Z
Core-2(config-if)# description TO_TX-ACCESS-2_PORT-22
Core-2(config-if)# vlan trunk allowed X11-X12
Core-2(config-if)#
```

21. Access ports **1/1/43 and 1/1/44**, make the port a trunk interface that allows **VLANs X11 and X12**.

```
Core-2(config-if)# exit
Core-2(config)# interface 1/1/43-1/1/44
Core-2(config-if-<1/1/43-1/1/44>)# vlan trunk allowed X11-X12
Core-2(config-if-<1/1/43-1/1/44>)# exit
```

```
Core-2(config)#
```

NOTICE: If you get the following error message:

```
Core-2(config-if-<1/1/43-1/1/44>)# vlan trunk allowed X11-X12
Operation not allowed on an interface part of a LAG (lag10).
```

This implies that your instructor has already run Lab - 6.1 - Link demonstration. This means that interface LAG 10 is replacing ports 43 and 44. Please configure the lag interface instead using this script:

```
Core-2(config)# interface lag 10
Core-2(config-lag-if)# vlan trunk allowed X11-X12
Core-2(config-lag-if)# exit
Core-2(config)#
```

Task 2: Verify the Topology

Objectives

Obtain and record the Bridge ID of the switches then, identify designated bridges for each link and locate the Root Bridge as well as link costs. This information will allow you to draw the current logical Common Spanning Tree (CST) topology.

Steps

Access-1

1. Access the terminal session to Access-1.
2. Show a filtered version of the "show spanning-tree" to get the switch MAC address and switch priority only.

```
T11-Access-1# show spanning-tree | begin 1 Bridge
  Bridge ID  Priority  : 32768
             MAC-Address: 88:3a:30:98:30:00
T11-Access-1#
```

IMPORTANT: Some of the command output depends on your switch hardware. For example, the system MAC address is unique to your equipment.

TIP: Since the output of the show spanning-tree command is quite long, we have decided to use a shorter version of it by displaying only the information that is relevant to us at this moment. You will use a regular version of this command later in this lab.

Was is the switch MAC address?

Was is the switch Priority?

3. Use this information to determine the Bridge ID of Access 1 and write down the value in **figure 5.1-2** down below.

TIP: You can obtain the Bridge ID by concatenating the value of Switch Priority with the Switch MAC address. e.g. 32768:**88:3a:30:98:30:00** for output above.

Core-1, Core-2 and Access-2

4. Repeat **steps 1 and 2** with **Core-1, Core-2 and Access-2**.

```
T11-Access-2# show spanning-tree | begin 1 Bridge
  Bridge ID  Priority  : 32768
             MAC-Address: 88:3a:30:97:a4:40
T11-Access-2#
```

```
Core-1# show spanning-tree | begin 1 Bridge
  Bridge ID  Priority  : 4096
             MAC-Address: 90:20:c2:bc:ed:00
Core-1#
```

```
Core-2# show spanning-tree | begin 1 Bridge
  Bridge ID  Priority  : 8192
             MAC-Address: 90:20:c2:bc:3f:00
Core-2#
```

5. On **figure 5.1-2** put a start by the switch that you have identified as **Root Bridge**. Other fields you will fill out in later steps.

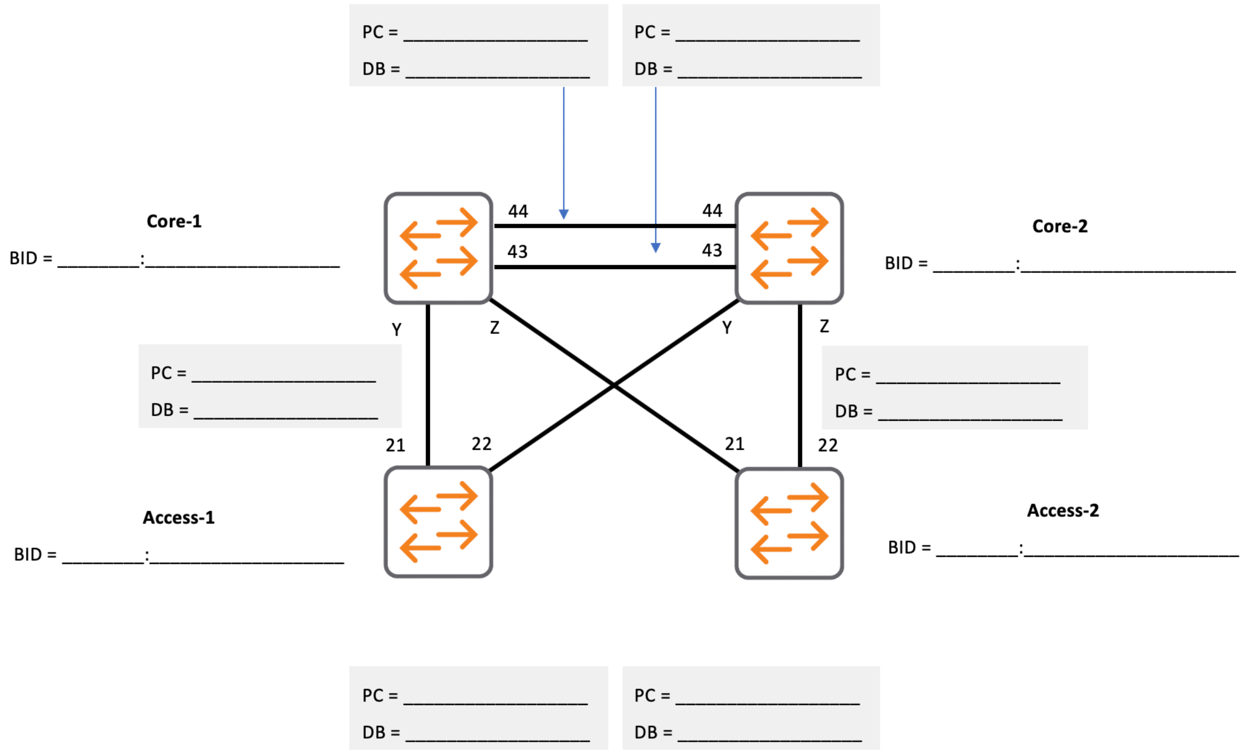

Figure 5.1-2: BIDs, Designated Bridges and costs.

TIP: You can find a larger copy of this diagram in Appendix 3.

Access-1

6. Move back to **Access-1** and run "**show spanning-tree**" command. What are the path costs of the ports?

```
T11-Access-1# show spanning-tree | begin 30 Port
Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Designated     Forwarding   20000          128        point_to_point
1/1/2        Disabled       Blocking     20000          128        point_to_point
1/1/3        Designated     Forwarding   20000          128        point_to_point
1/1/4        Disabled       Blocking     20000          128        point_to_point
1/1/5        Disabled       Blocking     20000          128        point_to_point
1/1/6        Disabled       Blocking     20000          128        point_to_point
                           ←---- output omitted ---→
```

What are path costs of ports?

7. All ports in this topology should have the same cost. Write down the path costs of all links on **figure 5.1-2**

IMPORTANT: Link path cost is relevant because it is used as a metric for calculating the Root Path Cost for each non-Root Bridge’s port. The port RPC is calculated by taking the RPC announcement in an incoming BDPU and adds it to the Link Path Cost of the port that receives the BPDU. This is equivalent to adding up the Link Patch Cost of each link between the local switch to the Root Bridge. If two or more ports have paths to the Root Bridge the one with the lowest Root Path Cost is the one that will be chosen as the Root Port.

RSTP (802.1r) and MST (802.1w) use path costs defined in the 802.1t standard which is an update of the legacy STP (802.1D). 802.1t defines the following path costs based on link speeds:

Link Speed	Value
100 Mbps	200,000
1 Gbps	20,000
10 Gbps	2,000
100 Gbps	200

8. Issue the “**show spanning-tree detail**” command. The output will be very long.

NOTE: “**show spanning-tree detail**” displays the role and state of the ports, similar to the “show spanning-tree” command, with the addition of which switch is the Designated Bridge for each link, the number of transitions to forwarding state, and the number of BPDUs being exchanged.

9. Now try now a filtered version of the “**show spanning-tree detail**” command in order to find the Designated bridge on each uplink.

```
T11-Access-1(config)# show spanning-tree  detail | begin 12 "Port 1/1/21"
Port 1/1/21 id 21
Designated root has priority              :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority            :4096 Address: 90:20:c2:bc:ed:00
Designated port id                        :16
Number of transitions to forwarding state  :11
```

```
Bpdus sent 1015, received 23535

Port 1/1/22 id 22
Designated root has priority               :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority             :8192 Address: 90:20:c2:bc:3f:00
Designated port id                         :37
Number of transitions to forwarding state  : 6
Bpdus sent 23, received 24183
T11-Access-1(config)#
```

What is the Switch's BID of the Designated Bridge on port 1/1/21 (port connected to Core-1)?

What is the designated port ID and who owns it?

What is the Switch's BID of the Designated Bridge on port 1/1/22 (port connected to Core-2)?

What is the designated port ID and who owns it?

10. Write down the Designated Bridge of these links on **figure 5.1-2**.

Access-2

11. Move to Access-2 and repeat **step 9**.

```
T11-Access-2# show spanning-tree detail | begin 12 "Port 1/1/21"
Port 1/1/21 id 21
Designated root has priority               :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority           :4096 Address: 90:20:c2:bc:ed:00
Designated port id                       :37
Number of transitions to forwarding state  : 1
Bpdus sent 5, received 21209

Port 1/1/22 id 22
Designated root has priority             :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority           :8192 Address: 90:20:c2:bc:3f:00
Designated port id                       :16
Number of transitions to forwarding state  : 1
Bpdus sent 7, received 21209
T11-Access-2#
```

What is the Switch’s BID of the Designated Bridge on port 1/1/21 (port connected to Core-1)?

What is the designated port ID and who owns it?

What is the Switch’s BID of the Designated Bridge on port 1/1/22 (port connected to Core-2)?

What is the designated port ID and who owns it?

12. Write down the Designated Bridge of these links on **figure 5.1-2**.

Core-2

13. Move to Core-2 and repeat **step 9** for ports **1/1/43 and 1/1/44**.

```
Core-2# show spanning-tree  detail | begin 12 "Port 1/1/43"
Port 1/1/43 id 43
Designated root has priority                 :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority               :4096 Address: 90:20:c2:bc:ed:00
Designated port id                           :43
Number of transitions to forwarding state    : 1
Bpdus sent 23556, received 23418

Port 1/1/44 id 44
Designated root has priority                 :4096 Address: 90:20:c2:bc:ed:00
Designated bridge has priority               :4096 Address: 90:20:c2:bc:ed:00
Designated port id                           :44
Number of transitions to forwarding state    : 2
Bpdus sent 23558, received 23416
Core-2#
```

What is the Switch's BID of the Designated Bridge on port 1/1/43 (port connected to Core-1)?

What is the designated port ID and who owns it?

What is the Switch's BID of the Designated Bridge on port 1/1/44?

What is the designated port ID and who owns it?

At this point you have obtained enough information to accurately determine the Root Bridge, the roles of ports from the Root Bridge to all the other switches, and to draw the CST topology. Start with the Root Bridge and ports' roles identification first.

Read the following notes in order to refresh how these elections take place.

IMPORTANT: Bridges role assignment are aligned with the following rules:
Rule 1: In a topology with redundant switch ports the Switch with lowest Bridge ID (Bridge Priority + MAC address) is elected Root Bridge.

Rule 2: A switch is considered to be closer to the Root Bridge if it has the lowest Root Path Cost from the root port and lowest BID combination. On a switch to switch link a designated bridge is the switch that is closest to the Root Bridge while the other switch will be non-designated bridge.

Rule 3: The Root Bridge is always the Designated bridge for all its links.

Rule 4: On a link connected to a collision domain where there is only one switch running STP, that switch will be the Designated Bridge for that link.

IMPORTANT: Port role assignment follows the following rules:

Rule 5: On a switch to switch link the port in the designated bridge side will be chosen as a designated port, unless there is a local loop on the same switch in which case the interface with the lowest Port ID will be designated port and the other will be the blocked port.

Rule 6: If a non-root bridge has only one switch to switch link, then the port used for that link is the Root Port.

Rule 7: If a non-root bridge has two or more switch-to-switch links to different remote devices, then:

a) The one with the lowest Root Path Cost is the root port. In case of a tie of two or more links with the same RPC then the one who's upstream switch is considered closest to the Root Bridge will be the Root port.
b) For any other links on which this switch was elected designated bridge, the interface will be chosen as designated port.

Rule 8: If a non-designated bridge has two or more links with equal RPC to the same Designated Bridge, then the local interface that connects neighbor's with lowest Port ID will be selected Root Port.

Rule 9: Any other interface on links where the local switch was not elected a designated bridge, will be considered an alternate port.

As a side note the final state of designated and root ports is Forwarding, unless there is a security feature triggering an action like root-guard, bpdu-protection, or loop-guard, in which case it will be either blocking or inconsistent.

Alternate ports final state will always be discarding.

Based on recorded information on figure 5.1-2, who is the Root Bridge? Remember that Root bridge is the switch with the lowest Bridge ID.

____

What was the Bridge ID component that made this switch the Root Bridge, the MAC address of the priority value?

____

____

Which switch will become Root if the current one fails?

____

14. Label the Root bridge **on figure 5.1-3**.
15. All Root Bridge's ports are Designated Ports tag them as DP on **figure 5.1-3**. → **Rule 3**.
16. Each Access Switch has two ports with different Root Path Costs (RPC), the one with the lowest value (20,000) is the root port (either port 21 or 22), tag them as RP. → **Rule 7a**.
17. The non-Root Core switch has two connections to the Root, since both have the same RPC value (20,000) the local port connected to the neighbor's interface with lowest Port ID will be the RP (interface 1/1/43) → **Rule 8**
18. On the other link between the non-Root Core Switch and Access-1, one of them is considered to be closest to the Root, that is the designated bridge, tag its port as DP. → **Rule 2, Rule 7b**.
19. Repeat **step 17** for the connection between the non-Root Bridge Core Switch and Access-2.
20. Last, both port Access Switches have one or two ports that are the only STP speaker (1/1/1 and 1/1/3 in Access-1 and 1/1/4 in Access-2). Therefore, Access Switches will be Designated Bridges for those ports, and the interfaces considered designated ports, tag them as DP → **Rule 4**
21. Any other interface will be considered an Alternate port. Draw an X on them to indicate the blocked link. → **Rule 9**

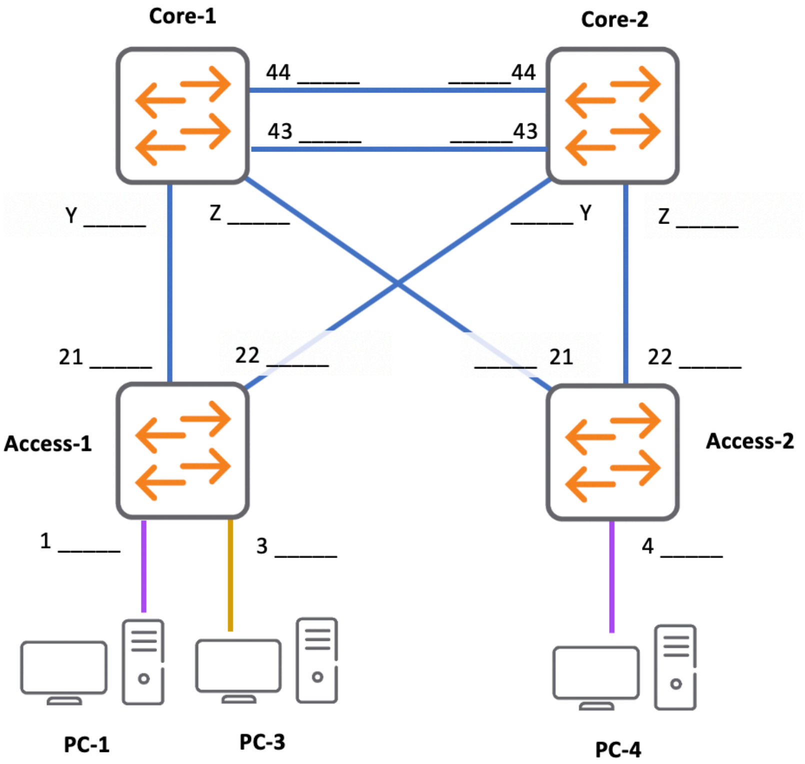

Figure 5.1-3: Devices and ports roles.

TIP: You can find a larger copy of this diagram in Appendix 3.

At this point you have a good idea of how the topology should look, in next steps this analysis will be validated.

22. On any switch run the filtered version of the "**show spanning-tree**" command (you should be currently on Core-2).

```
Core-2# show spanning-tree | begin 12 Spanning
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority   : 4096
             MAC-Address: 90:20:c2:bc:ed:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority  : 8192
             MAC-Address: 90:20:c2:bc:3f:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

Core-2#
```

What is the Bridge ID of the CST (MST0) Root Bridge?

Does the CST Root Bridge in the output match the one that you identified as in **figure 5.1-3**?

NOTE: The Root Bridge election result wasn’t random. By assigning low priority values of 4096 to Core-1 and 8192 to Core-2, Core-1 is elected root and Core-2 becomes the backup in case of failure. This is a best practice because at the Data Plane the Root acts as transport for traffic coming and going to devices connected to non-root bridges.

Core-1 and Core-2 (via PC-1)

23. Move to Core-1 and Core-2 then run the “**show running-config | include spanning-tree priority**”, and review the configuration used for manipulating the election.

```
Core-1# show running-config | include "spanning-tree priority"
spanning-tree priority 1
Core-1#
```

```
Core-2# show running-config | include "spanning-tree priority"
spanning-tree priority 2
Core-2#
```

IMPORTANT: 802.1D standard says that switch priority can be set in increments of 4096. AOS-CX reflects that rule by allowing the administrator to define a multiplying factor (called step) of this 4096 increment in a range between 0 and 15 where the default value is 8. See help output below:

```
Core-2(config)# spanning-tree priority ?
  <0-15>  Enter an integer number (Default: 8)
Core-2(config)# spanning-tree priority
```

Access-1 and Access-2

24. On Access Switches, use filtered versions of the “**show spanning-tree**” command for validating the roles of the ports.

```
T11-Access-1# show spanning-tree | exclude Disabled

                           ←---- output omitted ---→

Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Designated     Forwarding   20000          128        point_to_point
1/1/3        Designated     Forwarding   20000          128        point_to_point
1/1/21       Root           Forwarding   20000          128        point_to_point
1/1/22       Alternate      Blocking     20000          128        point_to_point
T11-Access-1#
```

```
T11-Access-2# show spanning-tree | exclude Disabled

                           ←---- output omitted ---→

Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Designated     Forwarding   20000          128        point_to_point
1/1/21       Root           Forwarding   20000          128        point_to_point
1/1/22       Alternate      Blocking     20000          128        point_to_point
T11-Access-2#
```

Do the outputs match your figure 5.1-3 results?

NOTE: If they do not, it may be because some of the ports are either down or the Access switches priorities are not 32768. Please fix that portion of the configuration before moving forward.

Core-1 and Core-2 (via PC-1)

25. On Core-1 and Core-2 use filtered versions of the "**show spanning-tree**" command for validating the roles of the ports. Look specifically for ports **1/1/Y, 1/1/Z, 1/1/43 and 1/1/44**.

NOTE: Replace the highlighted "Y", and "Z" for the port number of the downlinks that connect to Access-1 and Access-2.

```
Core-1# show spanning-tree | begin 40 Port
Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Designated     Forwarding   20000          128        point_to_point
1/1/2        Designated     Forwarding   20000          128        point_to_point
1/1/4        Designated     Forwarding   20000          128        point_to_point
1/1/5        Designated     Forwarding   20000          128        point_to_point
1/1/7        Designated     Forwarding   20000          128        point_to_point
1/1/8        Designated     Forwarding   20000          128        point_to_point
1/1/10       Designated     Forwarding   20000          128        point_to_point
1/1/11       Designated     Forwarding   20000          128        point_to_point
1/1/13       Designated     Forwarding   20000          128        point_to_point
1/1/14       Designated     Forwarding   20000          128        point_to_point
1/1/16       Designated     Forwarding   20000          128        point_to_point
1/1/17       Designated     Forwarding   20000          128        point_to_point
1/1/19       Designated     Forwarding   20000          128        point_to_point
1/1/20       Designated     Forwarding   20000          128        point_to_point
1/1/22       Designated     Forwarding   20000          128        point_to_point
1/1/23       Designated     Forwarding   20000          128        point_to_point
1/1/25       Designated     Forwarding   20000          128        point_to_point
1/1/26       Designated     Forwarding   20000          128        point_to_point
1/1/28       Designated     Forwarding   20000          128        point_to_point
1/1/29       Designated     Forwarding   20000          128        point_to_point
1/1/31       Designated     Forwarding   20000          128        point_to_point
1/1/32       Designated     Forwarding   20000          128        point_to_point
1/1/34       Designated     Forwarding   20000          128        point_to_point
1/1/35       Designated     Forwarding   20000          128        point_to_point
1/1/37       Designated     Forwarding   20000          128        point_to_point
1/1/38       Designated     Forwarding   20000          128        point_to_point
1/1/40       Designated     Forwarding   20000          128        point_to_point
1/1/41       Designated     Forwarding   20000          128        point_to_point
1/1/43       Designated     Forwarding   20000          128        point_to_point
```

```
1/1/44       Designated     Forwarding   20000          128        point_to_point
1/1/46       Designated     Forwarding   20000          128        point_to_point
1/1/47       Designated     Forwarding   20000          128        point_to_point
```

```
Core-2# show spanning-tree | begin 40 Port
Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/1        Designated     Forwarding   20000          128        point_to_point
1/1/2        Designated     Forwarding   20000          128        point_to_point
1/1/4        Designated     Forwarding   20000          128        point_to_point
1/1/5        Designated     Forwarding   20000          128        point_to_point
1/1/7        Designated     Forwarding   20000          128        point_to_point
1/1/8        Designated     Forwarding   20000          128        point_to_point
1/1/10       Designated     Forwarding   20000          128        point_to_point
1/1/11       Designated     Forwarding   20000          128        point_to_point
1/1/13       Designated     Forwarding   20000          128        point_to_point
1/1/14       Designated     Forwarding   20000          128        point_to_point
1/1/16       Designated     Forwarding   20000          128        point_to_point
1/1/17       Designated     Forwarding   20000          128        point_to_point
1/1/19       Designated     Forwarding   20000          128        point_to_point
1/1/20       Designated     Forwarding   20000          128        point_to_point
1/1/22       Designated     Forwarding   20000          128        point_to_point
1/1/23       Designated     Forwarding   20000          128        point_to_point
1/1/25       Designated     Forwarding   20000          128        point_to_point
1/1/26       Designated     Forwarding   20000          128        point_to_point
1/1/28       Designated     Forwarding   20000          128        point_to_point
1/1/29       Designated     Forwarding   20000          128        point_to_point
1/1/31       Designated     Forwarding   20000          128        point_to_point
1/1/32       Designated     Forwarding   20000          128        point_to_point
1/1/34       Designated     Forwarding   20000          128        point_to_point
1/1/35       Designated     Forwarding   20000          128        point_to_point
1/1/37       Designated     Forwarding   20000          128        point_to_point
1/1/38       Designated     Forwarding   20000          128        point_to_point
1/1/40       Designated     Forwarding   20000          128        point_to_point
1/1/41       Designated     Forwarding   20000          128        point_to_point
1/1/43       Root           Forwarding   20000          128        point_to_point
1/1/44       Alternate      Blocking     20000          128        point_to_point
1/1/46       Designated     Forwarding   20000          128        point_to_point
1/1/47       Designated     Forwarding   20000          128        point_to_point
Core-2#
```

Do the outputs match your **figure 5.1-3** results?

____

After validating your results, you are now ready to draw the CST which is the logical topology that will be used by switches for learning MAC addresses on each VLAN and determine how traffic is forwarded from all VLANs at Layer 2.

26. Based on your results and the current state of the diagram in figure 5.1-3, use **figure 5.1-4** to draw the CST. Use solid lines for active links and dotted lines for inactive ones.

NOTE: Active links are those with ports in forwarding mode at both sides of the cable while inactive links have an Alternate port on either side of the connection.

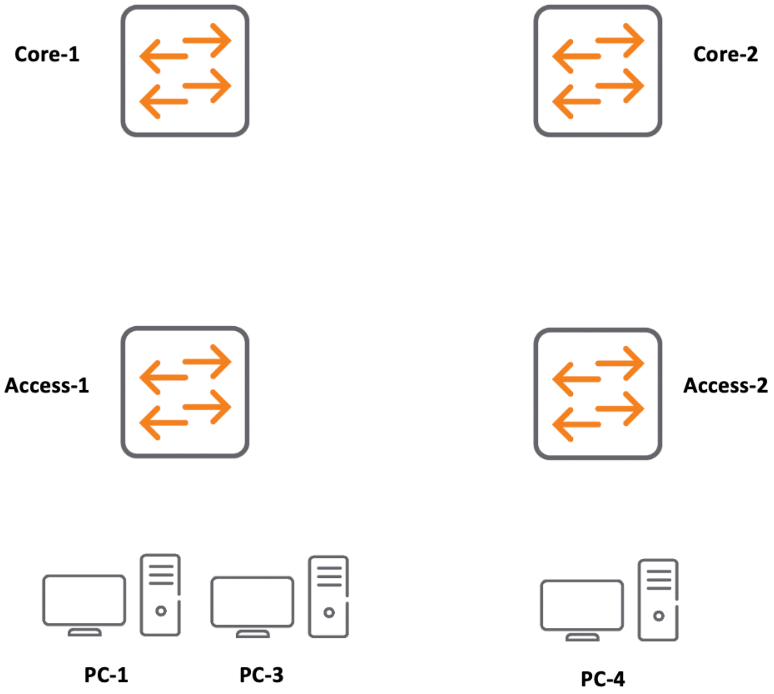

Figure 5.1-4: Drawing CST

TIP: You can find a larger copy of this diagram in Appendix 3.

Resultant topology should be similar to figure 5.1-5:

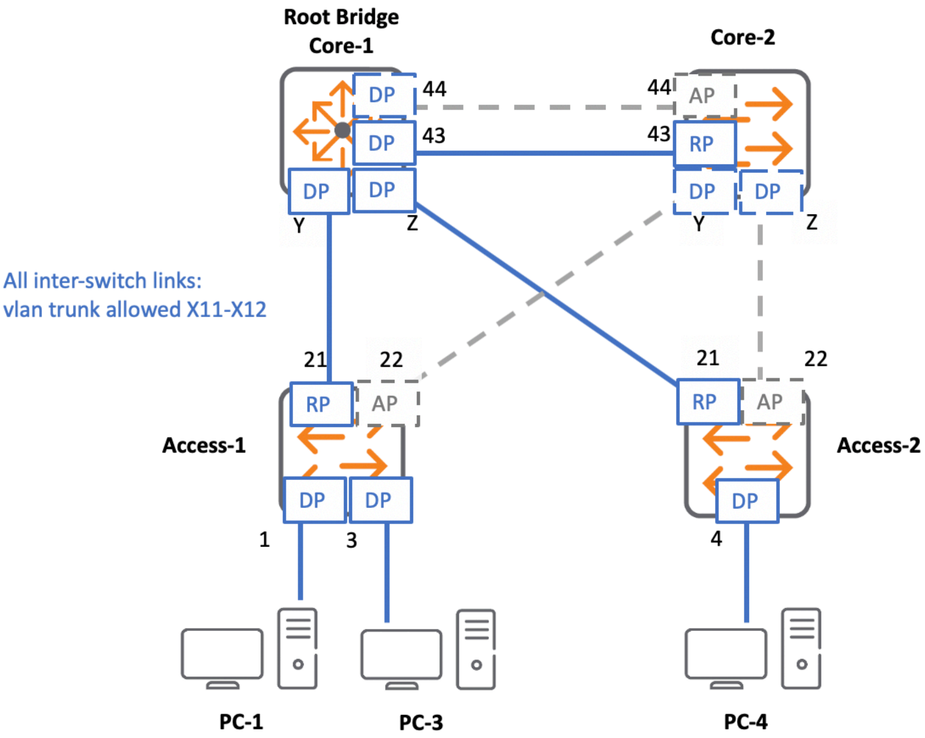

Figure 5.1-5: CST Topology detailed

Task 3: Test Link Failure.

Objectives

After discovering the CST topology, you should have a good idea of how traffic flows, you will now test how resilient the network is to a failure of any uplink.

Steps

PC-1

1. Access PC-1 and run a continuous ping to PC-4 (**10.X.12.104**). Ping should be successful.

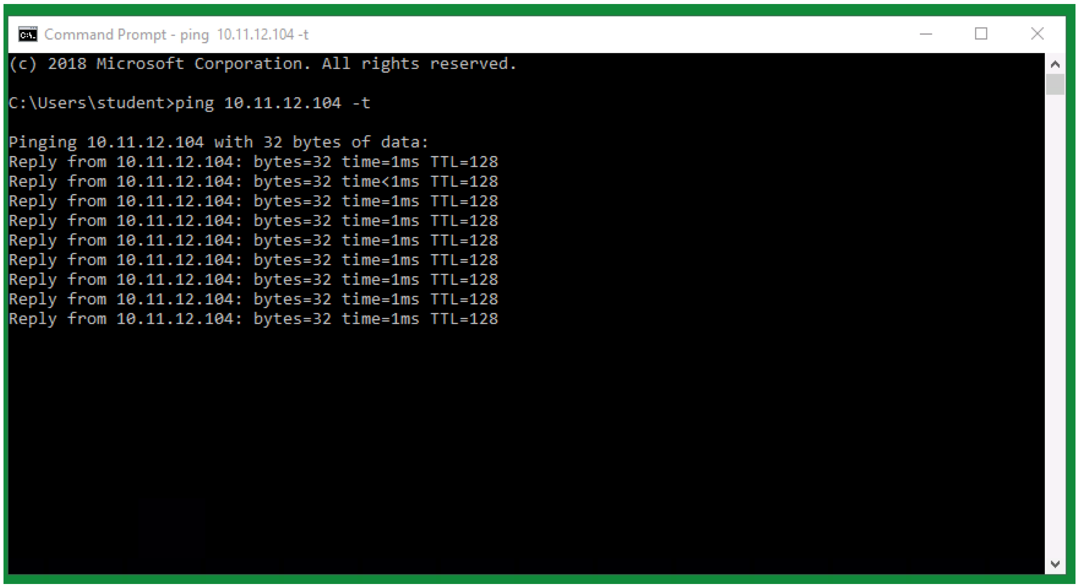

```
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.12.104 -t

Pinging 10.11.12.104 with 32 bytes of data:
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time<1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
Reply from 10.11.12.104: bytes=32 time=1ms TTL=128
```

Figure 5.1-6: Ping to PC-4.

IMPORTANT: At this point and based on figure 5.1-3 traffic is flowing from PC-1 to Access-1 → Access-1 to Core-1 (using port 1/1/21 to 1/1/Y link) → Core-1 to Access-2 (using 1/1/Z to 1/1/21 link) → Access-2 to PC-4. You will now modify the topology and analyze the traffic path.

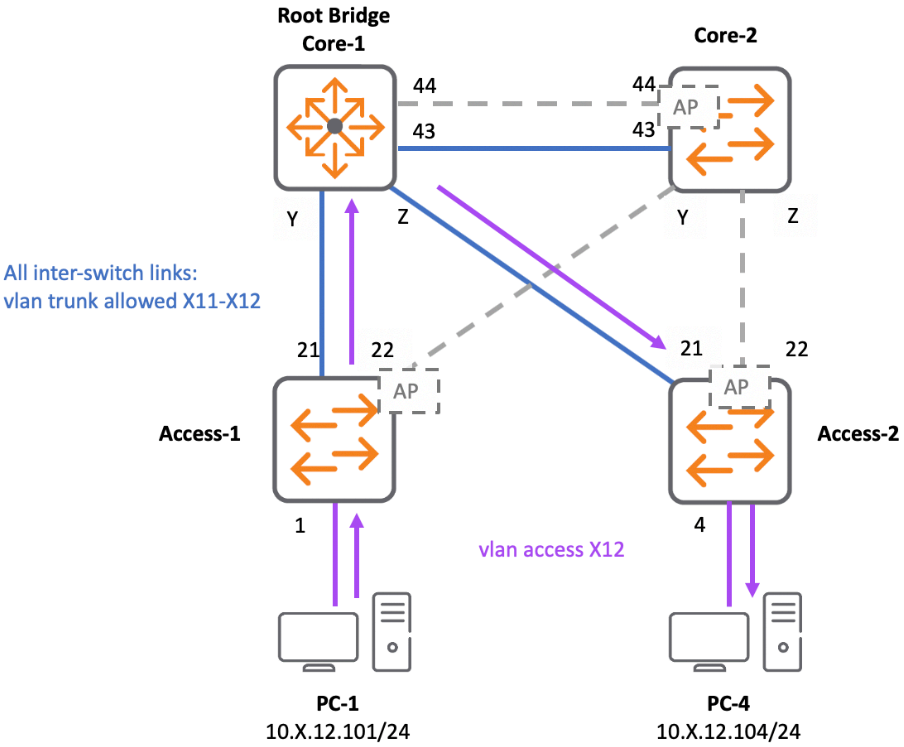

Figure 5.1-7: CST Topology

Access-1

2. Move to Access-1 and use the "**show spanning-tree**" command to verify the current Root port. It should be 1/1/21.

```
T11-Access-1# show spanning-tree | include Root
   Root ID.  Priority    : 4096
1/1/21       Root              Forwarding    20000           128         point_to_point
T11-Access-1#
```

3. Disable port 1/1/21.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/21
T11-Access-1(config-if)# shutdown
T11-Access-1(config-if)#
```

4. Repeat **step 2**.

```
T11-Access-1(config-if)# show spanning-tree | include Root
   Root ID.  Priority   : 4096
1/1/22       Root           Forwarding   20000          128        point_to_point
T11-Access-1#
```

PC-3

5. Move back to PC-1 and verify the ping.

Is ping still running?

____

How many packets did you lose?

____

What is the traffic flow now?

____

IMPORTANT: Traffic is now flowing from PC-1 to Access-1 → Access-1 to Core-2 (using port 1/1/22 to 1/1/Y link) → Core-2 to Core-1 using port 1/1/43 link → Core-1 to Access-2 (using port 1/1/Z to 1/1/21 link) → Access-2 to PC-4. As seen in figure 5.1-8.

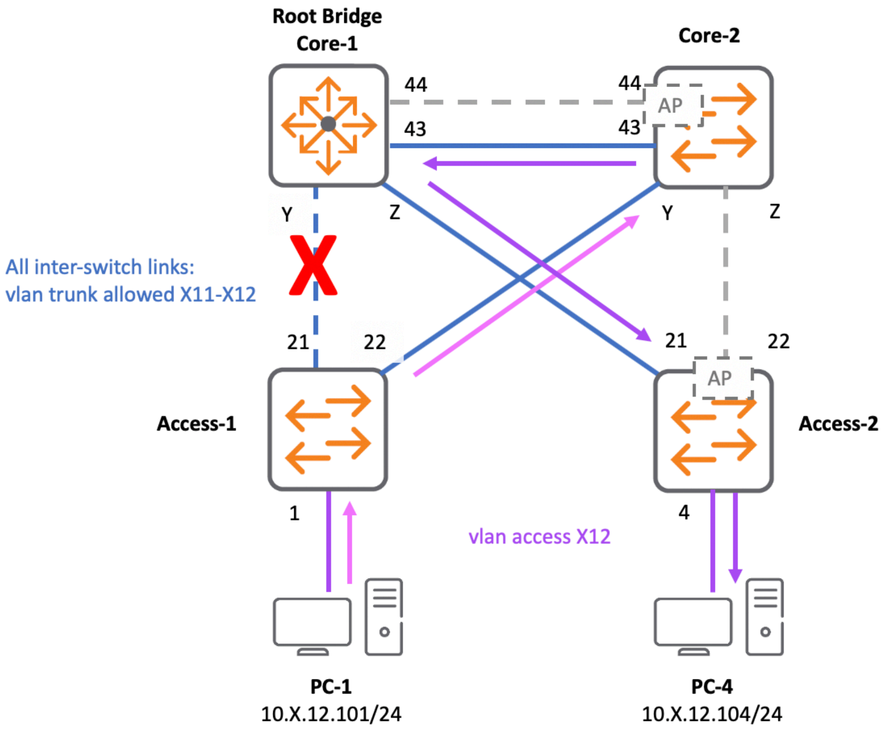

Figure 5.1-8: CST Topology after failure.

Access-1

6. Move to **Access-1** and re-enable port 1/1/21. The topology should return to normal.

```
T11-Access-1(config-if)# no shutdown
T11-Access-1(config-if)#
```

```
T11-Access-1(config-if)# show spanning-tree | include Root
   Root ID.  Priority   : 4096
1/1/21       Root          Forwarding   20000          128        point_to_point
T11-Access-1(config-if)#
```

Task 4: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

Access-1, Access-2, Core-1 and Core-2 (via PC-1)

1. Save the current Access and Core switches' configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

Access-1 and Access-2.

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab5-1_final**.

```
T11-Access-1# copy running-config checkpoint Lab5-1_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab5-1_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 5.1!

AOS-CX Switching Fundamentals

Lab 5.2: Deploying MSTP

Overview

Surprisingly enough, two days after the second Core was deployed a fiber connection was broken in the MDF. This affected the Access-1 main uplink, however your previous STP configuration avoided any network disruption. BigStartup (your customer) only realized there was a failure in the link when they received notification from Rent4Cheap Properties. Your customer is very satisfied with your advice. Your business relationship and their trust in you is growing.

Nonetheless, the failover event made BigStartup management wonder: Are the uplinks in an idle state when there is no failure? Are there connections that normally do not forward any traffic? Is it possible to share the load across those uplinks?

When you were asked those questions, the answer was "yes" to all of them. You went on to explain there is a new version of the STP protocol that not only provides loop avoidance and fast failover, but also provides load sharing and that it could be easily deployed. It is called Multiple Instance Spanning Tree. The next morning you received a request to deploy the solution.

Objectives

After completing this lab, you will be able to:

- Deploy an MST Region Configuration
- Draw per instance topologies
- Validate the load sharing effect

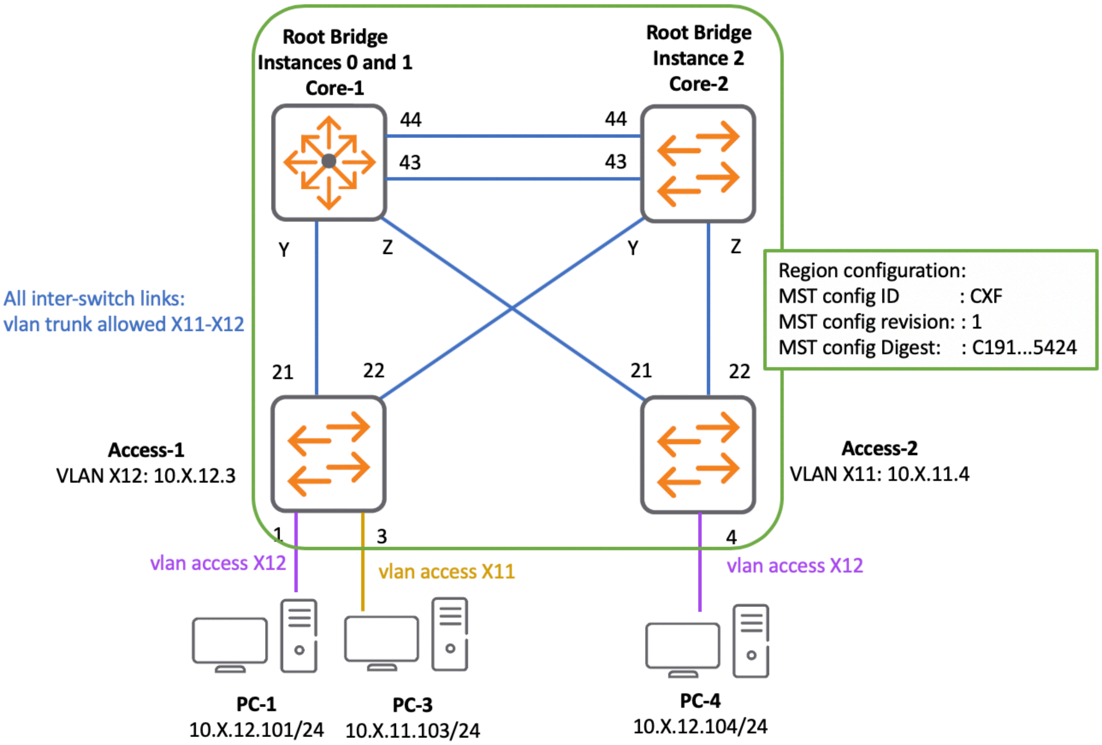

Figure 5.2-1: Lab Topology

Task 1: Inspect MST Region Configuration

Objectives

Core switches have been pre-provisioned with an MST region configuration that cannot be modified. Therefore, in this lab you will deploy the same MST region script on your Access Switches. Then you will explore the current Core's priority values and confirm that all switches agree on the Root Bridge in each Instance.

Core-1 (via PC-1)

1. Access Core-1.
2. Display the current MST region configuration.

```
Core-1(config)# show spanning-tree mst-config
MST configuration information
   MST config ID        : CXF
   MST config revision  : 1
   MST config digest    : C1918786A14CE2765D013B62CCCD5424
   Number of instances  : 2

Instance ID     Member VLANs
--------------- ----------------------------------
0               1-110,113-210,213-310,313-410,413-510,513-610,613-710,713-810,813-
910,913-1010,
               1013-1110,1113-1210,1213-1310,1313-1410,1413-4094
1               111,211,311,411,511,611,711,811,911,1011,1111,1211,1311,1411
2               112,212,312,412,512,612,712,812,912,1012,1112,1212,1312,1412

Core-1(config)#
```

What are the MST config ID and revision number values?

What is the config digest value?

What is the Instance to VLAN mapping configuration?

Instance 1: ____

____

Instance 2: ____

____

NOTE: Since the Core switches are a shared resource in a multitenancy environment, several VLANs terminate on them. Although many of these VLANs are not applicable to your environment, they must be part of the MST Region configuration in order to distribute these VLANs' traffic across multiple uplinks based on the Root Bridge of each instance.

IMPORTANT: The MST config digest is the result of hashing the instance to VLAN mapping configuration. The digest along with the region ID (region name) and revision number are contained within the MST BDPUs sent by the switches. Switches transmit their region to one another. If the region announced in an incoming BPDU matches the local MST configuration, then the local switch forms part of its neighbor's region. Switches belonging to the same region converge towards each instance's root bridge and form part of each instance's topology.

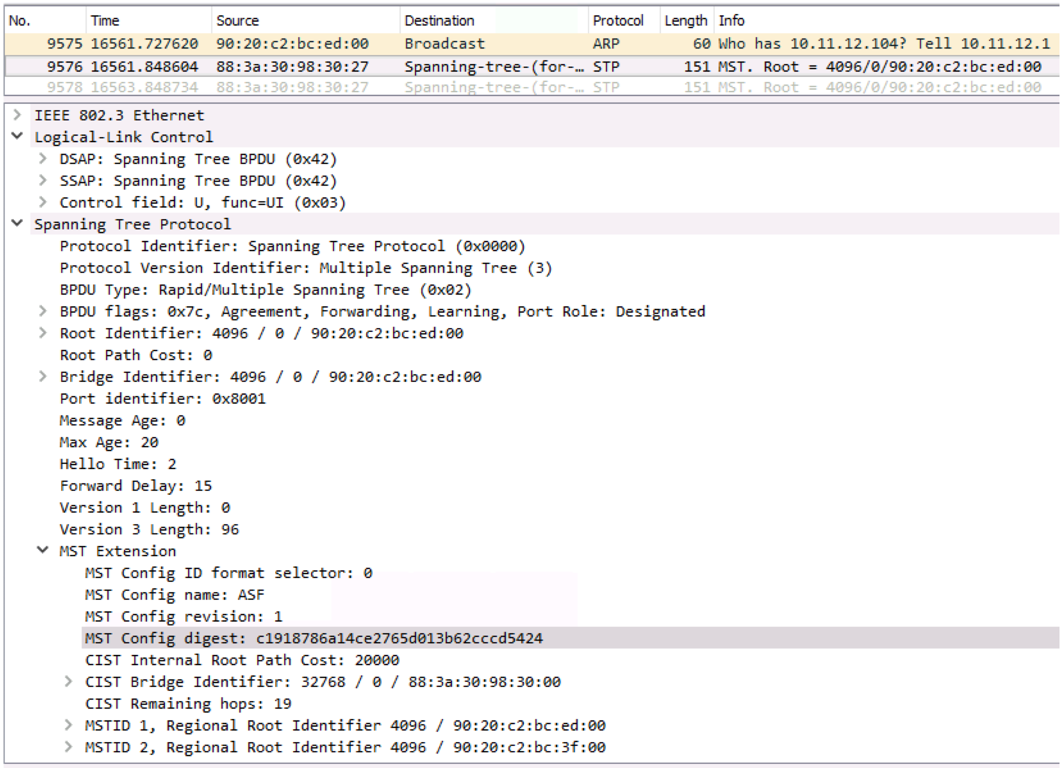

```
No.   Time               Source              Destination            Protocol  Length  Info
9575  16561.727620       90:20:c2:bc:ed:00   Broadcast              ARP       60      Who has 10.11.12.104? Tell 10.11.12.1
9576  16561.848604       88:3a:30:98:30:27   Spanning-tree-(for-…   STP       151     MST. Root = 4096/0/90:20:c2:bc:ed:00
9578  16563.848734       88:3a:30:98:30:27   Spanning-tree-(for-…   STP       151     MST. Root = 4096/0/90:20:c2:bc:ed:00
```

```
> IEEE 802.3 Ethernet
v Logical-Link Control
  > DSAP: Spanning Tree BPDU (0x42)
  > SSAP: Spanning Tree BPDU (0x42)
  > Control field: U, func=UI (0x03)
v Spanning Tree Protocol
    Protocol Identifier: Spanning Tree Protocol (0x0000)
    Protocol Version Identifier: Multiple Spanning Tree (3)
    BPDU Type: Rapid/Multiple Spanning Tree (0x02)
  > BPDU flags: 0x7c, Agreement, Forwarding, Learning, Port Role: Designated
  > Root Identifier: 4096 / 0 / 90:20:c2:bc:ed:00
    Root Path Cost: 0
  > Bridge Identifier: 4096 / 0 / 90:20:c2:bc:ed:00
    Port identifier: 0x8001
    Message Age: 0
    Max Age: 20
    Hello Time: 2
    Forward Delay: 15
    Version 1 Length: 0
    Version 3 Length: 96
  v MST Extension
      MST Config ID format selector: 0
      MST Config name: ASF
      MST Config revision: 1
      MST Config digest: c1918786a14ce2765d013b62cccd5424
      CIST Internal Root Path Cost: 20000
    > CIST Bridge Identifier: 32768 / 0 / 88:3a:30:98:30:00
      CIST Remaining hops: 19
    > MSTID 1, Regional Root Identifier 4096 / 90:20:c2:bc:ed:00
    > MSTID 2, Regional Root Identifier 4096 / 90:20:c2:bc:3f:00
```

Figure 5.2-2: STP BPDU with MST extensions

Core-2 (via PC-1)

3. Move to Core-2 and repeat step 2.

```
Core-2(config)# show spanning-tree mst-config
MST configuration information
   MST config ID          : CXF
   MST config revision    : 1
   MST config digest      : C1918786A14CE2765D013B62CCCD5424
   Number of instances    : 2

Instance ID     Member VLANs
--------------- ----------------------------------
0               1-110,113-210,213-310,313-410,413-510,513-610,613-710,713-810,813-
910,913-1010,
               1013-1110,1113-1210,1213-1310,1313-1410,1413-4094
1               111,211,311,411,511,611,711,811,911,1011,1111,1211,1311,1411
2               112,212,312,412,512,612,712,812,912,1012,1112,1212,1312,1412

Core-2(config)#
```

Do region parameters match the ones of Core-1?

ANSWER: It does, this confirms that both Core switches are part of the same region, however your Access switches are not since they do not have any custom region configuration.

Access-1

4. Move to Access-1 and use the "**show spanning-tree**" command. Then move to Access-2 and use it again.

```
T11-Access-1# show spanning-tree mst-config
MST configuration information
   MST config ID        : 88:3a:30:98:30:00
   MST config revision  : 0
   MST config digest    : AC36177F50283CD4B83821D8AB26DE62
   Number of instances  : 0

Instance ID     Member VLANs
--------------- ----------------------------------
0               1-4094

T11-Access-1#
```

What is the default Config ID and revision number?

What is the default VLAN to Instance mapping?

Access-2

5. Move to Access-2 and use a filtered version of the same command.

```
T11-Access-2# show spanning-tree mst-config | include MST
MST configuration information
   MST config ID        : 88:3a:30:97:a4:40
   MST config revision  : 0
   MST config digest    : AC36177F50283CD4B83821D8AB26DE62
T11-Access-2#
```

Are Access switches in the same region as the Core switches?

Are the two Access switches part of the same region?

ANSWER: As you can see, the Access switches' configuration is different from the Core switches and although Access-1 and Access-2 share the same Digest (result of having all VLANs mapped to Instance 0) they do not share the region ID or revision number and therefore they belong to different regions. See figure 5.2-3.

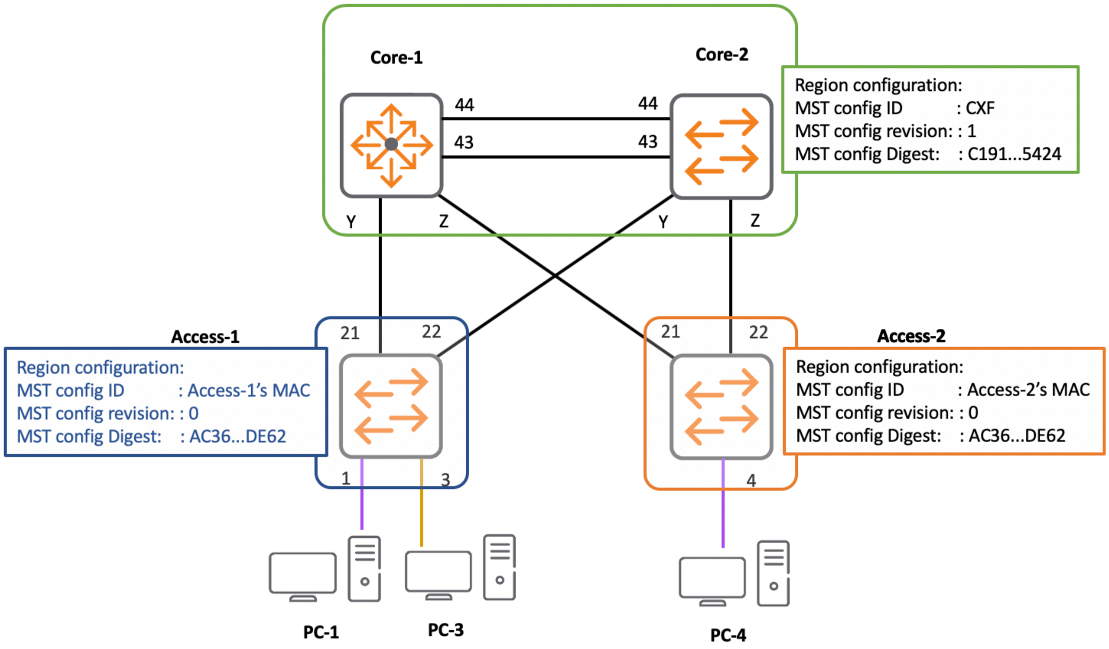

Region configuration:
MST config ID : CXF
MST config revision: : 1
MST config Digest: : C191...5424

Region configuration:
MST config ID : Access-1's MAC
MST config revision: : 0
MST config Digest: : AC36...DE62

Region configuration:
MST config ID : Access-2's MAC
MST config revision: : 0
MST config Digest: : AC36...DE62

Figure 5.2-3: Multi-Region topology.

IMPORTANT: Switches that do not share a common region configuration will belong to different regions, if this is the case then they will run RSTP, negotiate roles within the CST and form part of the CST topology only. They will lack any MST based load sharing support. In this type of design, root and designated ports will forward traffic for all VLANs and similarly alternate ports will discard traffic from all VLANs.

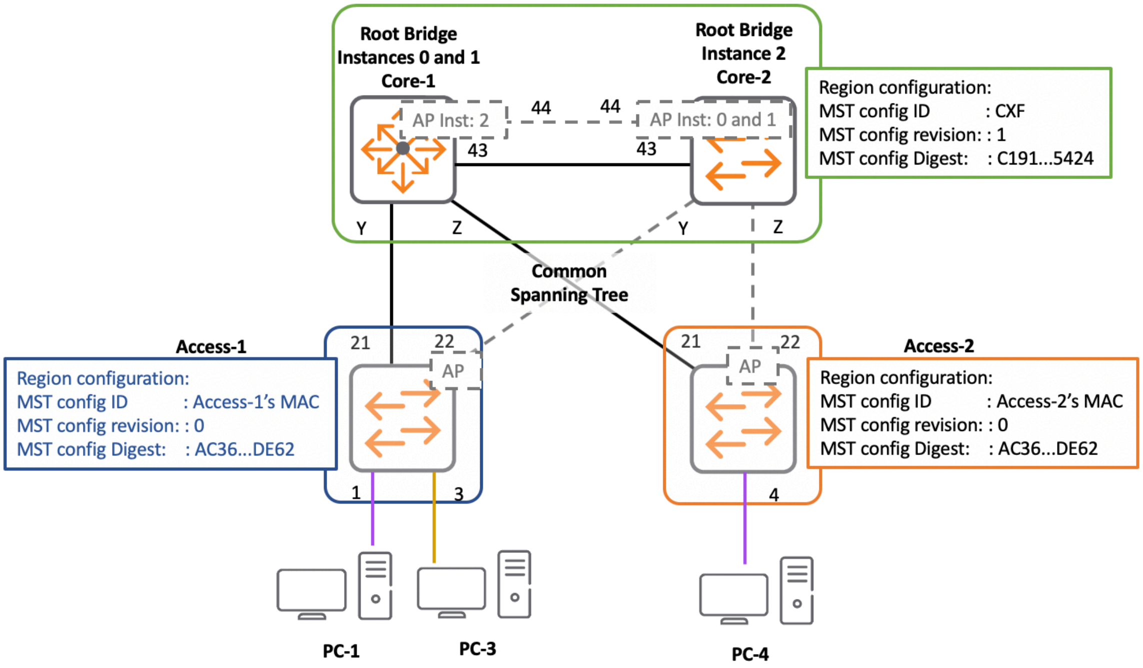

Figure 5.2-4: Multi region CST.

Task 2: Inspect Load Balancing

Objectives

Confirm what link Access-1 is using for each VLAN by inspecting its MAC Address table, then apply the same Core switch configuration to the Access switches and inspect the MAC table.

This test is easy for VLAN X12, because PC-1 and PC-4 (members of that VLAN) are connected to different access switches and their traffic has to cross the Core. However, testing VLAN X11 is more difficult because there is a single client (PC-3) on Access-1. In order to generate IP traffic on VLAN X11, you will simulate a host on Access-2 by adding an IP address on that switch using Switched Virtual Interfaces (SVI).

Steps

Access-2

1. Move to Access-2's console.
2. Create **interface vlan X11**, then assign it IP address **10.X.11.4/24**

> **NOTE:** Replace the highlighted "X" for your student table number.

```
T11-Access-2(config)#
T11-Access-2(config)# interface vlan X11
T11-Access-2(config-if-vlan)# ip address 10.X.11.4/24
T11-Access-2(config-if-vlan)# exit
```

3. See the newly created SVI details using "**show ip interface vlanX11**"

```
T11-Access-2# show ip interface vlanX11

Interface vlan1111 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 88:3a:30:97:a4:40
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IPv4 address 10.11.11.3/24
 L3 Counters: Rx Disabled, Tx Disabled

Access-2#
```

IMPORTANT: This command is case sensitive, make sure to type lowercase “vlan” immediately followed by the VLAN number, e.g. “**show ip interface vlanX11**”

What is the SVI state?

____

4. Record the MAC address of Interface VLAN X11 of Access-2.

Access-2’s MAC address ____

PC-4

5. Access PC-4.
6. Record the MAC address of PC-4.

PC-4’s MAC address ____

TIP: You can also record the MAC address values on figure 5.1-4 of the handouts.

PC-3

7. Access PC-3.
8. Run a continuous ping to Access-2 IP address on **VLAN X11** (**10.X.11.4**). Ping should be successful.

PC-1

9. Access PC-1.
10. Run a continuous ping to PC-4's IP address on **VLAN X12** (**10.X.12.104**). Ping should be successful.

Access-1.

11. Move back to Access-1.
12. Display the MAC address table.

```
T11-Access-1(config)# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                          Port
--------------------------------------------------------------
88:3a:30:97:a4:40    1111     dynamic                       1/1/21
00:50:56:b1:30:69    1111     dynamic                       1/1/1
00:50:56:b1:ae:e8    1111     dynamic                       1/1/3
00:50:56:b1:a9:86    1112     dynamic                       1/1/21
T11-Access-1(config)#
```

What port is used to reach Access-2's MAC address?

What port is used to reach PC-4's MAC address?

13. Apply the STP region configuration:
    - Config-name: **CXF**
    - Config-revision: **1**
    - Instance 1 VLANs: **111, 211, 311, 411, 511, 611, 711, 811, 911, 1011, 1111, 1211, 1311 and 1411**.
    - Instance 2 VLANs: **112, 212, 312, 412, 512, 612, 712, 812, 912, 1012, 1112, 1212, 1312 and 1412**.

TIP: You can find a copy of that script in the “CXF Region Configuration.txt” contained in the “CXF Student Folder” on PC-1’s desktop. Feel free to copy and paste it when needed.

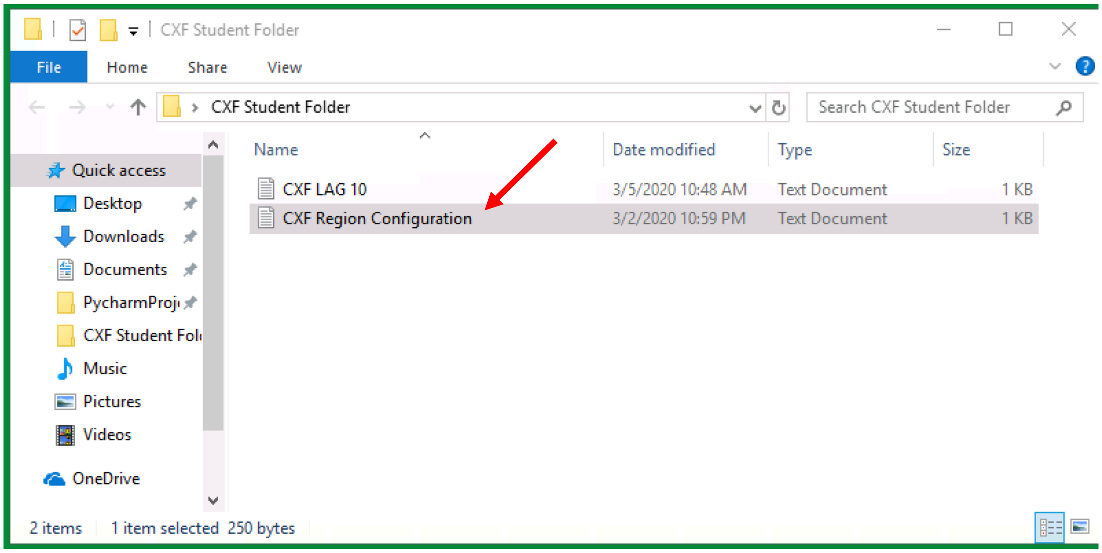

```
T11-Access-1# configure terminal
T11-Access-1(config)# spanning-tree config-name CXF
T11-Access-1(config)# spanning-tree config-revision 1
T11-Access-1(config)# spanning-tree instance 1 vlan
111,211,311,411,511,611,711,811,911,1011,1111,1211,1311,1411
T11-Access-1(config)#spanning-tree instance 2 vlan
112,212,312,412,512,612,712,812,912,1012,1112,1212,1312,1412
T11-Access-1(config)# end
```

NOTICE: You should be careful when applying the region configuration. The smallest difference will make the integration into the region fail. Config-name is case sensitive, revision level must be “1” in this case, and every single VLAN listed in the script must be included regardless of whether they apply to your table or not.

14. Confirm config ID, revision number and digest match the ones seen on Task 1 step 3.

```
T11-Access-1# show spanning-tree mst-config | include MST
MST configuration information
   MST config ID        : CXF
   MST config revision  : 1
   MST config digest    : C1918786A14CE2765D013B62CCCD5424
T11-Access-1#
```

15. Move to Access-2 and repeat steps 12 and 13.

```
T11-Access-2# configure terminal
T11-Access-2(config)# spanning-tree config-name CXF
T11-Access-2(config)# spanning-tree config-revision 1
T11-Access-2(config)# spanning-tree instance 1 vlan
111,211,311,411,511,611,711,811,911,1011,1111,1211,1311,1411
T11-Access-2(config)# spanning-tree instance 2 vlan
112,212,312,412,512,612,712,812,912,1012,1112,1212,1312,1412
T11-Access-1(config)# end
```

```
T11-Access-2# show spanning-tree mst-config | include MST
MST configuration information
   MST config ID        : CXF
   MST config revision  : 1
   MST config digest    : C1918786A14CE2765D013B62CCCD5424
T11-Access-2#
```

NOTE: At this point Spanning Tree is running 3 processes simultaneously, one per instance. The topology that is used is 100% dependent on who the Root is for each instance, which in turn depends on the BID of the switches. Currently Access switches have no custom priority whatsoever, but the Cores are already provisioned with certain values, please proceed and validate those values.

Core-1 (via PC-1)

16. Move to Core-1 and explore its STP priorities configuration.

```
Core-1(config)# show running-config | include priority
spanning-tree priority 1
spanning-tree instance 1 priority 1
spanning-tree instance 2 priority 2
Core-1(config)#
```

Core-2 (via PC-1)

17. Move to Core-2 and repeat the previous step.

```
Core-2(config)# show running-config | include priority
spanning-tree priority 2
spanning-tree instance 1 priority 2
spanning-tree instance 2 priority 1
Core-2(config)#
```

Based on the outputs, who is the Root for each instance?

Root for Instance 0:____

Root for Instance 1:____

Root for Instance 2:____

IMPORTANT: Instance 0 or Internal Spanning Tree (IST) is used as both: a regular instance in MST and the creation of the CST in a multi-region deployment for backward compatibility with RSTP speakers, for this reason Instance 0 is known as CIST (Common Internal Spanning Tree)

Validate your conclusions.

Access-1

18. Move to Access-1.
19. Use the "**show spanning-tree mst 0**" command to look at information about instance 0.

TIP: Since this command's output is long, a filtered version of it is used below.

```
T11-Access-1# show spanning-tree mst 0 | begin 25 Root | exclude Disabled
Root           Address:90:20:c2:bc:ed:00  Priority:4096
               Port:1/1/21                Path cost:0
Regional Root  Address:90:20:c2:bc:ed:00  Priority:4096
               Internal cost:20000        Rem Hops:19

Port           Role           State        Cost       Priority   Type
```

```
-------------- -------------- ------------ ---------- ---------- ----------
1/1/1          Designated     Forwarding   20000      128        point_to_point
1/1/3          Designated     Forwarding   20000      128        point_to_point
1/1/21         Root           Forwarding   20000      128        point_to_point
1/1/22         Alternate      Blocking     20000      128        point_to_point

Topology change flag          : True
Number of topology changes    : 9
Last topology change occurred : 1422 seconds ago

T11-Access-1#
```

Who is the Root bridge for this instance?

____

What are the Root and Alternate ports?

____

20. Repeat step 19 for instances 1 and 2.

```
T11-Access-1# show spanning-tree mst 1 | begin 25 Root | exclude Disabled
Root           Address:90:20:c2:bc:ed:00    Priority:4096
               Port:1/1/21, Cost:20000, Rem Hops:19

Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
1/1/1          Designated     Forwarding   20000   128        point_to_point
1/1/3          Designated     Forwarding   20000   128        point_to_point
1/1/21         Root           Forwarding   20000   128        point_to_point
1/1/22         Alternate      Blocking     20000   128        point_to_point

Topology change flag          : True
Number of topology changes    : 4
Last topology change occurred : 1449 seconds ago

T11-Access-1#
```

Who is the regional root for this instance?

____

What are the Root and Alternate ports?

```
T11-Access-1# show spanning-tree mst 2 | begin 25 Root | exclude Disabled
Root            Address:90:20:c2:bc:3f:00    Priority:4096
                Port:1/1/22, Cost:20000, Rem Hops:19

Port            Role            State         Cost    Priority   Type
--------------  --------------  ------------  ------- ---------- ----------
1/1/1           Designated      Forwarding    20000   128        point_to_point
1/1/3           Designated      Forwarding    20000   128        point_to_point
1/1/21          Alternate       Blocking      20000   128        point_to_point
1/1/22          Root            Forwarding    20000   128        point_to_point

Topology change flag          : True
Number of topology changes    : 5
Last topology change occurred : 1488 seconds ago

T11-Access-1#
```

Who is the regional root for this instance?

What are the Root and Alternate ports?

TIP: There is no need to validate the same information on Access-2. Since it has the same region configuration, the results will be the same.

NOTE: As you can see Instance 0 and 1 share the same Root and same roles on uplinks, however Instance 2 does not. This is because Core-2 is root for this instance. Instance topologies are similar to the ones in figures below.

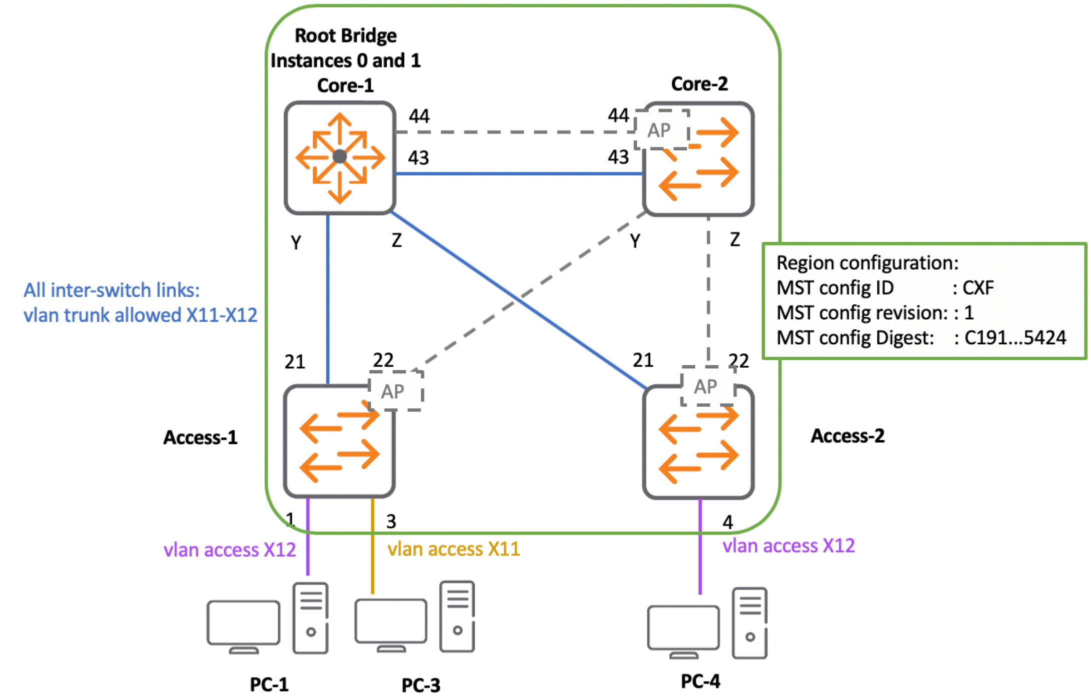

Figure 5.2-5: Topology of instances 0 and 1

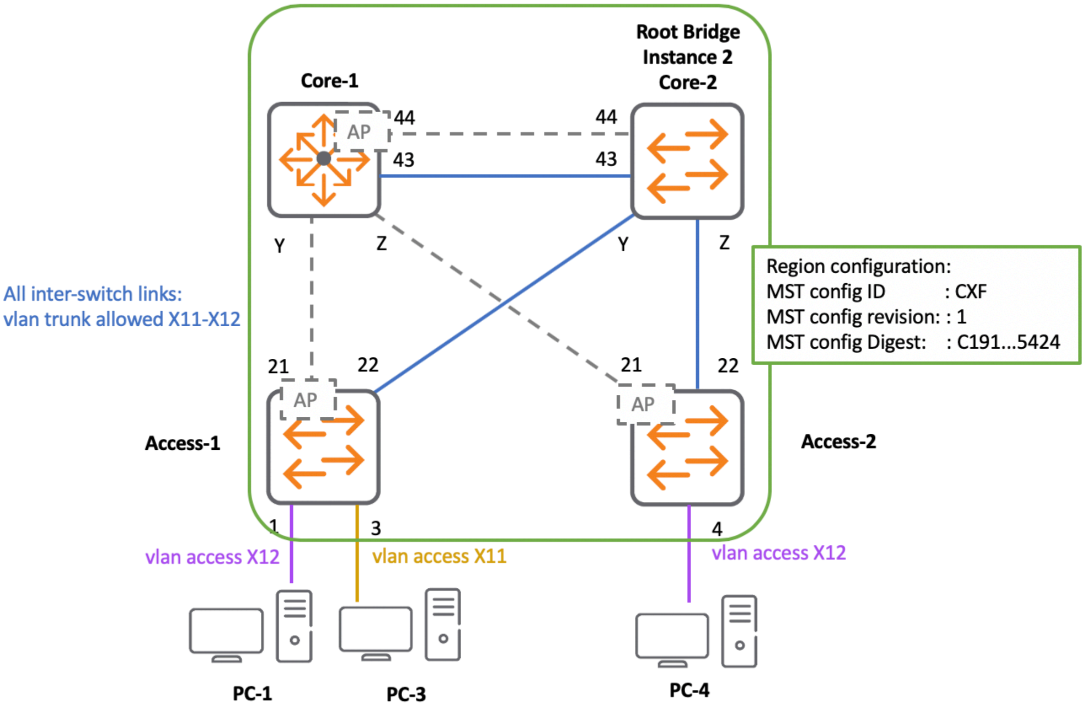

Figure 5.2-6: Topology of instance 2.

Finally, you will inspect the MAC address table, if everything is correct the MAC address of PC-4 should be seen now on a different port.

21. Display the MAC address table.

```
T11-Access-1(config)# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                      Port
--------------------------------------------------------------
88:3a:30:97:a4:40    1111     dynamic                   1/1/21
00:50:56:b1:30:69    1111     dynamic                   1/1/1
00:50:56:b1:ae:e8    1111     dynamic                   1/1/3
00:50:56:b1:a9:86    1112     dynamic                   1/1/22
T11-Access-1(config)#
```

What port is used to reach Access-2’s MAC address?

____

What port is used to reach PC-4’s MAC address?

____

What has changed from what you saw in step 7?

____

PC-1 and PC-3

22. Stop the pings.

Task 3: Save Your Configurations

Objectives

Save your configurations and create checkpoints. Note that lab checkpoints might be used by later activities.

Steps

Access-1 and Access-2.

1. Save the current Access and Core switches' configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab5-2_final**.

```
T11-Access-1# copy running-config checkpoint Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 5.2!

AOS-CX Switching Fundamentals

Lab 5.3: Explore Broadcast Storm Effects (optional)

Overview

In a previous Module you were introduced to the potential problems that a layer 2 loop can bring. In this lab activity you will intentionally create one by creating a dual-home topology between the two access switches and removing spanning-tree. Also, you will use two alternative methods for containing and preventing such loops that can be used in addition to Spanning-tree.

Objectives

After completing this lab, you will be able to:

- Create a redundant topology
- Force a layer 2 loop and create a broadcast and multicast storm
- Find evidence of the layer two loop
- Prevent loops using loop-protect

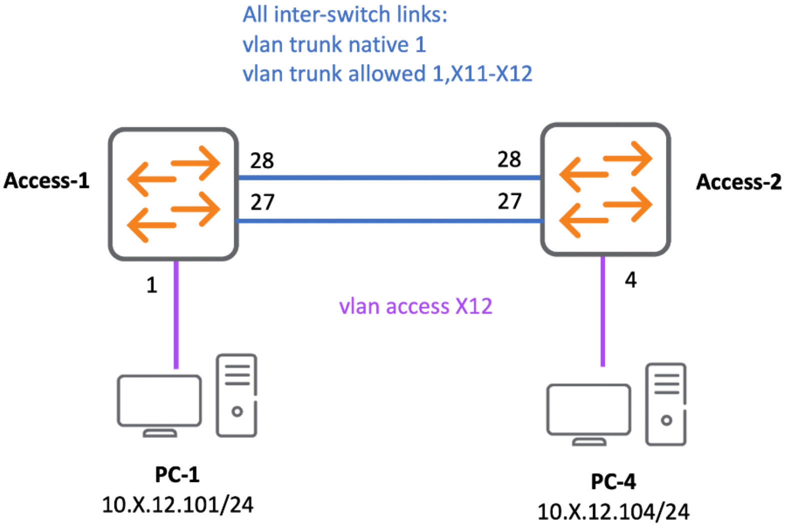

Figure 5-3-1: Lab Topology

Task 1: Pre-lab Setup:

Objectives

In this activity you will isolate Access-1 and Access-2 from the rest of the network then enable a dual homed topology using ports 27 and 28.

Steps

Access-1

1. Open a console connection to Access-1. Login using **admin** and **no password**.
2. Disable ports **1/1/21 and 1/1/22**.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/21-1/1/22
T11-Access-1(config-if-<1/1/21-1/1/22>)# shutdown
T11-Access-1(config-if-<1/1/21-1/1/22>)# exit
```

3. Access port **1/1/27** and set a description.

```
T11-Access-1(config)# interface 1/1/27
T11-Access-1(config-if)# description TO_TX-ACCESS-2_PORT-27
T11-Access-1(config-if)# exit
```

4. Create a port range including **1/1/27 and 1/1/28**, allow **VLANs 1, X11 and X12**, then enable them.

```
T11-Access-1(config)# interface 1/1/27-1/1/28
T11-Access-1(config-if-<1/1/27-1/1/28>)# vlan trunk allow 1,X11-X12
T11-Access-1(config-if-<1/1/27-1/1/28>)# no shutdown
T11-Access-1(config-if-<1/1/27-1/1/28>)# exit
```

5. Confirm port **1/1/21 to 1/1/22** are down.

```
T11-Access-1(config) # show interface 1/1/21-1/1/22 link-status
--------------------------------------------------
```

```
Port      Type            Physical   Link
                          Link State Transitions
--------------------------------------------------
1/1/21    1GbT            down       4
1/1/22    1GbT            down       2
T11-Access-1(config) # exit
```

NOTICE: Remember that you are about to create a Layer 2 loop, which has the potential of affecting other students. In order to limit the effects, you have to make sure that both uplinks 1/1/21 and 1/1/22 are down. **If this is not the case, go to those ports and shut them down.**

Access-2

6. Move to **Access-2**, then repeat **steps 2 through 4**.

```
T11-Access-2# configure terminal
T11-Access-2(config-if-<1/1/21-1/1/22>)# interface 1/1/21-1/1/22
T11-Access-2(config-if-<1/1/21-1/1/22>)# shutdown
T11-Access-2(config-if-<1/1/21-1/1/22>)# exit
```

```
T11-Access-2(config)# interface 1/1/27
T11-Access-2(config-if)# description TO_TX-ACCESS-1_PORT-27
T11-Access-2(config-if)# exit
```

```
T11-Access-2(config)# interface 1/1/27-1/1/28
T11-Access-2(config-if-<1/1/27-1/1/28>)# vlan trunk allow 1,X11-X12
T11-Access-2(config-if-<1/1/27-1/1/28>)# no shutdown
T11-Access-2(config-if-<1/1/27-1/1/28>)# exit
```

7. Confirm ports **1/1/21 to 1/1/22** are down and **1/1/27 to 1/1/28** are up.

```
T11-Access-2(config)# show interface 1/1/21-1/1/22,1/1/27-1/1/28 link-status
--------------------------------------------------
Port      Type            Physical   Link
                          Link State Transitions
--------------------------------------------------
1/1/21    1GbT            down       0
1/1/22    1GbT            down       4
1/1/27    SFP+DAC1        up         11
1/1/28    SFP+DAC1        up         17
T11-Access-2(config)#
```

NOTICE: Remember that you are about to create a Layer 2 loop, which has the potential of affecting the entire network, in order to limit the effects, you have to make sure that both uplinks 1/1/21 and 1/1/22 are down. **Do not proceed if this is not the case.**

8. Increase Access-2 spanning tree priority to 15 (61440). This will make Access-1 the root bridge and force Access-2 to choose a root and alternate port.

```
T11-Access-2(config)# spanning-tree priority 15
T11-Access-2(config)# exit
```

9. Use the show spanning-tree command and look at **1/1/27 and 1/1/28** ports.

```
T11-Access-2# show spanning-tree | exclude Disabled
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority   : 28672
             MAC-Address: 88:3a:30:98:30:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority  : 61440
             MAC-Address: 88:3a:30:97:a4:40
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

Port         Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/4        Designated     Forwarding   20000          128        point_to_point
1/1/27       Root           Forwarding   2000           128        point_to_point
1/1/28       Alternate      Blocking     2000           128        point_to_point
```

What interface is the root port?

____

Was interface is the alternate port?

____

Since the current Access-1 and Access-2 configurations will be used later, create checkpoints now.

Access-1 and Access-2

10. Backup the current Access switches’ configuration as a custom checkpoint called **Lab5-3_task1_done**.

```
T11-Access-1# copy running-config checkpoint Lab5-3_task1_done
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab5-3_task1_done
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

Task 2: Create Layer 2 Loop.

Objectives

In this task you will create a layer 2 loop and a broadcast/multicast storm as a consequence of that. Then you will witness the symptoms and gather logs that document its presence.

Steps

Access-1

1. Open a console connection to **Access-1**.
2. Clear 1/1/27 and 1/1/28 interfaces' statistics. Then display those interfaces statistics.

```
T11-Access-1# clear interface 1/1/27 statistics
T11-Access-1# clear interface 1/1/28 statistics
T11-Access-1# show interface 1/1/27-1/1/28 statistics
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
Interface                         RX Bytes          RX Packets             TX Bytes
TX Packets         RX Broadcast          RX Multicast         TX Broadcast         TX
Multicast
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
1/1/27                                 369                   3                  476
2                  0                     3                   1                    1
1/1/28                                 123                   1                    0
0                  0                     1                   0                    0
T11-Access-1#
```

In total, how many broadcast and multicast packets has 1/1/27 received since the count was last cleared?

____

In total, how many broadcast and multicast packets has 1/1/28 received since the count was last cleared?

____

3. Wait a minute then repeat **step 2**.

```
T11-Access-1# show interface 1/1/27-1/1/28 statistics
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
Interface                        RX Bytes           RX Packets             TX Bytes
TX Packets        RX Broadcast         RX Multicast         TX Broadcast         TX
Multicast
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
1/1/27                               1852                   15                  822
3                   0                  15                    2                   1
1/1/28                               1952                   14                    0
0                   1                  13                    0                   0
T11-Access-1#
```

How many total broadcast and multicast packets do you see on both interfaces?

____

Have the number of packets statistics grown by dozens, hundreds or thousands?

____

PC-4

4. Access PC-4 and issue the "**ipconfig -all**" command and record PC-4's MAC address.

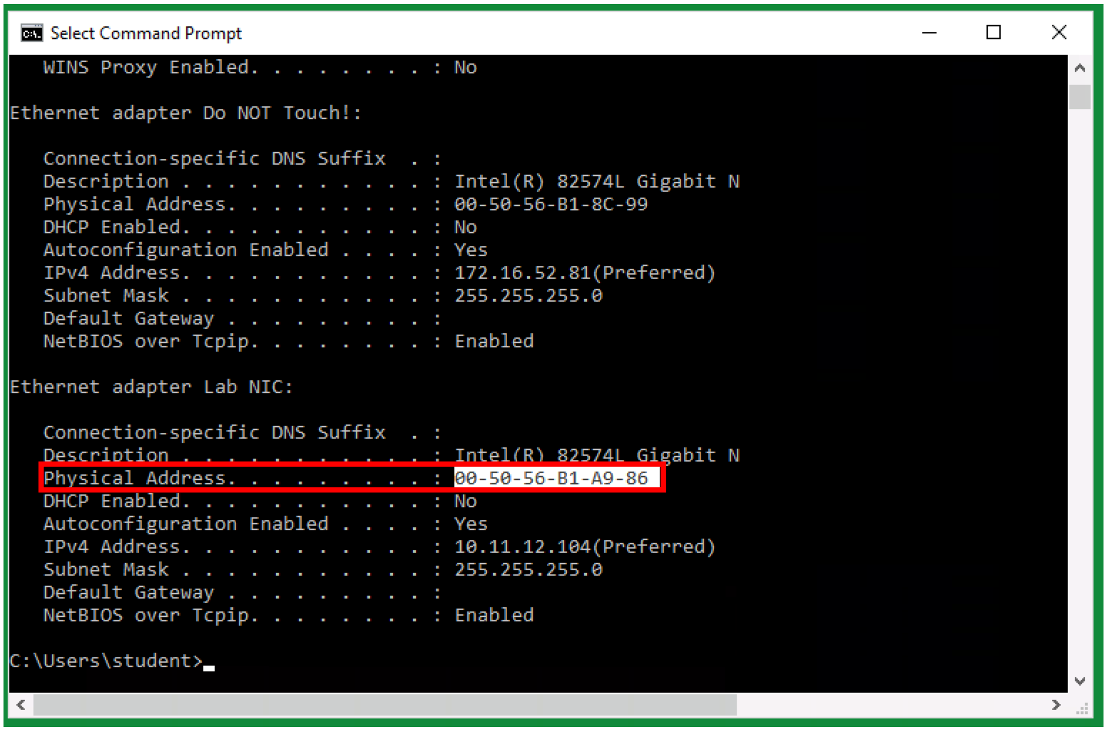

```
   WINS Proxy Enabled. . . . . . . . : No

Ethernet adapter Do NOT Touch!:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit N
   Physical Address. . . . . . . . . : 00-50-56-B1-8C-99
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 172.16.52.81(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   NetBIOS over Tcpip. . . . . . . . : Enabled

Ethernet adapter Lab NIC:

   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit N
   Physical Address. . . . . . . . . : 00-50-56-B1-A9-86
   DHCP Enabled. . . . . . . . . . . : No
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 10.11.12.104(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
   NetBIOS over Tcpip. . . . . . . . : Enabled

C:\Users\student>_
```

Figure 5.3-2: PC-4's MAC Address

5. Run a continuous ping to PC-1's IP address (**10.X.12.101**). Ping should be successful.

Access-2

6. Move back to Access-2.
7. Enable Spanning-Tree BPDU-filtering on interfaces **1/1/27 and 1/1/28.**

NOTE: BPDU-filtering is a feature that prevents a switch from sending or receiving Spanning Tree BPDUs. When enabling the feature on ports 1/1/27 and 1/1/28, you will prevent Access-2 from processing incoming Access-1's BPDUs and also Access-1 will no longer receive Access-2's BPDUs. This will cause, after few seconds, a transition.

NOTE: Connecting a device with BPDU filtering enabled to an Access Switch in order to create a layer 2 loop is a well-known Denial of Service attack. Later, in task 3, you will learn an effective way of protecting your network against this threat.

```
T11-Access-2(config)# interface 1/1/27-1/1/28
T11-Access-2(config-if-<1/1/27-1/1/28>)# spanning-tree bpdu-filter
T11-Access-2(config-if-<1/1/27-1/1/28>)# end
```

8. Use the show spanning-tree command and look at current 1/1/27 and 1/1/28 interfaces state in Access-2. They will now be in Forwarding mode.

```
T11-Access-2# show spanning-tree | include Forwarding
1/1/4        Designated     Forwarding   20000          128        point_to_point
1/1/27       Designated     Forwarding   2000           128        point_to_point
1/1/28       Designated     Forwarding   2000           128        point_to_point
T11-Access-2#
```

Does this create a Layer 2 loop?

____

ANSWER: There should now be a loop and a broadcast storm. You will now gather evidence of its presence.

Access-1

9. Move back to Access-1
10. Wait a minute and display the Access-1 interfaces statistics again.

```
T11-Access-1# show interface 1/1/27-1/1/28 statistics
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
Interface                        RX Bytes           RX Packets             TX Bytes
TX Packets         RX Broadcast         RX Multicast         TX Broadcast          TX
Multicast
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
------------------
1/1/27                           99708452               447079            100828088
449949                  0                446704                    2             449761
```

```
1/1/28                              121948168          546263          123196578
549809                    1              545890                0          549405
```

How many more broadcast and multicast packets combined do you have now on each interface?

IMPORTANT: The large increment of broadcast and multicast packet in a short period of time is one piece of evidence of a broadcast storm. It is the result of a loop.

PC-4

11. Move back to PC4 and look at the connectivity test.

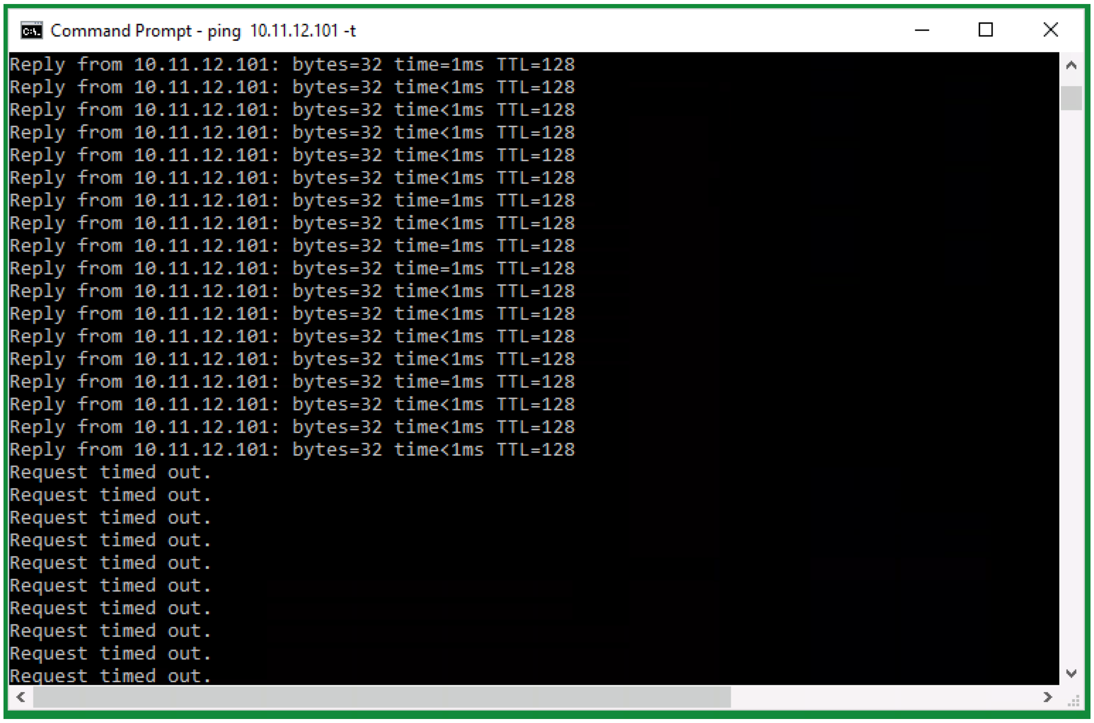

```
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time=1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Reply from 10.11.12.101: bytes=32 time<1ms TTL=128
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

Figure 5-3-3: Ping failure.

Is the ping working flawlessly?

IMPORTANT: The lack of connectivity in the affected devices is one of the main symptoms of a broadcast storm.

12. Stop the ping.

Access-1

13. Move to **Access-1**
14. Enable layer-2 mac event debugs (l2mac event debugs), set the buffer as the debug destination, then enable paging.

```
T11-Access-1# debug l2mac event
T11-Access-1# debug destination buffer
T11-Access-1# page
```

15. Show the debug buffer of the L2MAC module using the “include” filtering command followed by the 4 last hexadecimal characters of PC1’s MAC address that you recorded in step 2.

```
T11-Access-1# show debug buffer module L2MAC | include a9:86
2020-01-14:16:59:53.240738|l2mac-
mgrd|LOG_DEBUG|MSTR|1|L2MAC|L2MAC_EVENT|macmgr_mac_manager_handle_mac_event(1311)
, MAC=00:50:56:b1:a9:86, VLAN=1112, Port=1/1/27 is trying to be inserted
2020-01-14:16:59:53.240785|l2mac-
mgrd|LOG_DEBUG|MSTR|1|L2MAC|L2MAC_EVENT|macmgr_mac_manager_handle_mac_event(1328)
, MAC=00:50:56:b1:a9:86 was successfully inserted
2020-01-14:16:59:54.241041|l2mac-
mgrd|LOG_DEBUG|MSTR|1|L2MAC|L2MAC_EVENT|macmgr_mac_manager_handle_mac_event(1311)
, MAC=00:50:56:b1:a9:86, VLAN=1112, Port=1/1/28 is trying to be inserted
2020-01-14:16:59:54.241089|l2mac-
mgrd|LOG_DEBUG|MSTR|1|L2MAC|L2MAC_EVENT|macmgr_mac_manager_handle_mac_event(1328)
, MAC=00:50:56:b1:a9:86 was successfully inserted
                                            ←---- output omitted ---→
T11-Access-1#q
```

NOTICE: If your MAC address includes letters as part of the hexadecimal notation, then make sure to type them in lower case as in example above: “**a**9:86”.

Are there any events describing mac address learning on interface 1/1/27 first, then 1/1/28 later?

IMPORTANT: A MAC address learning of flapping between all interfaces involved in the loop is another piece of evidence of a broadcast storm. The affected interfaces are not necessarily the ones where the client is connected!!!

16. Display the system information.

```
T11-Access-1# show system
Hostname           : Access-1
System Description : FL.10.04.0030
System Contact     :
System Location    :

Vendor             : Aruba
Product Name       : JL668A 6300F 24G 4SFP56 Sw
Chassis Serial Nbr : SG90KN70HX
Base MAC Address   : 883a30-983000
AOS-CX Version : FL.10.04.0003

Time Zone          : UTC

Up Time            : 5 days, 6 hours, 16 minutes
CPU Util (%)       : 49
Memory Usage (%)   : 17
T11-Access-1#
```

What is the current CPU utilization?

TIP: If CPU increase is not that evident, then you can also try the same verification command on Access-2.

If you remember from Lab 2, average utilization was always below 10%

IMPORTANT: The final indication of a broadcast storm is high CPU utilization.

Task 3: Contain a Broadcast Storm.

Objectives

In this task you will enable a port-based feature called Rate Filtering that controls the number of Broadcast and Multicast packets per second. It is important to know that the layer 2 loop will still be present, but we are considerably attenuating its effects.

Steps

Access-1

1. Access a port range that includes ports **1/1/27 and 1/1/28**.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/27-1/1/28
```

2. Enable rate filtering for broadcast and multicast traffic allowing a maximum of 75 packets per second.

```
T11-Access-1(config-if-<1/1/27-1/1/28>)# rate-limit broadcast 75 pps
T11-Access-1(config-if-<1/1/27-1/1/28>)# rate-limit multicast 75 pps
T11-Access-1(config-if-<1/1/27-1/1/28>)#
```

3. Display the system CPU utilization.

```
T11-Access-1(config-if-<1/1/27-1/1/28>)# show system | include CPU
CPU Util (%)       : 3
T11-Access-1# Access-1(config-if-<1/1/27-1/1/28>)#
```

What is the current CPU utilization now?

____

IMPORTANT: Rate limit is a good protection mechanism when connecting devices to your network that you don’t control.

Next you will test another feature that might be of use for preventing loops. To test its affects you will have to remove rate filtering in order to re-create the unstable situation once more.

4. Remove rate-limit related commands.

```
T11-Access-1(config-if-<1/1/27-1/1/28>)# no rate-limit broadcast
T11-Access-1(config-if-<1/1/27-1/1/28>)# no rate-limit multicast
T11-Access-1(config-if-<1/1/27-1/1/28>)#
```

5. Wait a few seconds then verify the system CPU, it should have risen once more.

```
T11-Access-1(config-if-<1/1/27-1/1/28>)# show system | include CPU
CPU Util (%)       : 45
T11-Access-1(config-if-<1/1/27-1/1/28>)# exit
```

Task 4: Preventing Loops.

Objectives

In this task you will deploy loop-protect which blocks ports involved in loops. Although this feature is intended to be in place before the loop happens, applying it now will demonstrate its ability to bring the network back to normal.

Steps

Access-1

1. Move to Access-1's console.
2. Enable loop protect in ports **1/1/27 and 1/1/28**, then set tx-rx-disable as the loop-protect action. The feature will take effect immediately.

```
T11-Access-1(config)# interface 1/1/27-1/1/28
T11-Access-1(config-if-<1/1/27-1/1/28>)# loop-protect
T11-Access-1(config-if-<1/1/27-1/1/28>)# loop-protect action tx-rx-disable
T11-Access-1(config-if-<1/1/27-1/1/28>)# end
```

3. Display interfaces where loop-protect has detected loops.

```
T11-Access-1# show loop-protect loop-detected

Status and Counters - Loop Protection Information

Transmit Interval              : 5 (sec)
Port Re-enable Timer           : Disabled

Interface 1/1/27
  Loop-protect enabled         : Yes
  Action on loop detection     : TX RX disable
  Loop detected count          : 1
  Loop detected                : Yes
    Detected on VLAN           : 1
    Detected at                : 2020-01-14T17:02:20
  Interface status             : down

Interface 1/1/28
  Loop-protect enabled         : Yes
  Action on loop detection     : TX RX disable
  Loop detected count          : 1
  Loop detected                : Yes
    Detected on VLAN           : 1
    Detected at                : 2020-01-14T17:02:20
  Interface status             : down

T11-Access-1#
```

What interfaces are listed in the output?

____

Have loops been detected on them?

____

What are the interface status now?

____

NOTE: In the example above the switch was able to detect the loop on both ports simultaneously, however it is also possible to have the switch detecting the loop on one of the ports first and blocking it before detecting it on the other.

4. Use the "**show interface brief**" command for displaying the current state of ports 1/1/27 and 1/1/28.

```
T11-Access-1# show interface brief | exclude Administratively
--------------------------------------------------------------------------------
----
Port      Native  Mode   Type           Enabled Status  Reason                Speed
          VLAN                                                               (Mb/s)
--------------------------------------------------------------------------------
----
1/1/1     1112    access 1GbT             yes     up                           1000
1/1/3     1111    access 1GbT             yes     up                           1000
1/1/27    1       trunk  SFP+DAC1         yes     down     Network loop detected  --
1/1/28    1       trunk  SFP+DAC1         yes     down     Network loop detected  --
vlan1     --             --               yes     down                           --
T11-Access-1#
```

Are ports 1/1/27 and 1/1/28 administratively enabled?

What is the status of the port?

What is the reason behind this status?

5. Display the system CPU utilization, you will see how the value has normalized again.

```
T11-Access-1# show system | include CPU
CPU Util (%)       : 3
T11-Access-1#
```

You will now proceed to remove the loop.

Access-2

6. Move to Access-2.
7. Remove BPDU filtering from ports 1/1/27 and 1/1/28.

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface 1/1/27-1/1/28
T11-Access-2(config-if-<1/1/27-1/1/28>)# no spanning-tree bpdu-filter
T11-Access-2(config-if-<1/1/27-1/1/28>)# end
```

Access-1

8. Move to Access-1.
9. Remove loop-protect from ports 1/1/27 and 1/1/28.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/27-1/1/28
T11-Access-1(config-if-<1/1/27-1/1/28>)# no loop-protect
T11-Access-1(config-if-<1/1/27-1/1/28>)# no loop-protect  action
```

10. Disable ports 1/1/27 and 1/1/28, then enable them back. This will remove the “Network loop detected” state and bring them back on.

```
T11-Access-1(config-if-<1/1/27-1/1/28>)# shutdown
T11-Access-1(config-if-<1/1/27-1/1/28>)# no shutdown
T11-Access-1(config-if-<1/1/27-1/1/28>)# end
```

Task 5: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

Access-1 and Access-2.

1. Save the current Access switches' configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab5-3_final**.

```
T11-Access-1# copy running-config checkpoint Lab5-3_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab5-3_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 5.3!

AOS-CX Switching Fundamentals

Lab 5.4: Fine tune MSTP (Instructor Demo - Optional)

Overview

Your customer is pleased with the current MST setup, especially after showing the load balancing effect across the Access switches uplinks. Now he wonders if that benefit is also present on links between both Cores. When BigStartup brings that up you not only get surprised of how much they are getting involved with the projects but you also realize that load sharing isn't available there because regardless the instance, link on ports 43 is always active while 44 is inactive.

In this lab you will fine tune MSTP settings in order to use both Core to Core links.

Objectives

After completing this lab, you will be able to:

- Inspect default port priorities.
- Manipulate port priorities on Core-1.
- Change the Root port on Core-2 for instances 0 and 1.

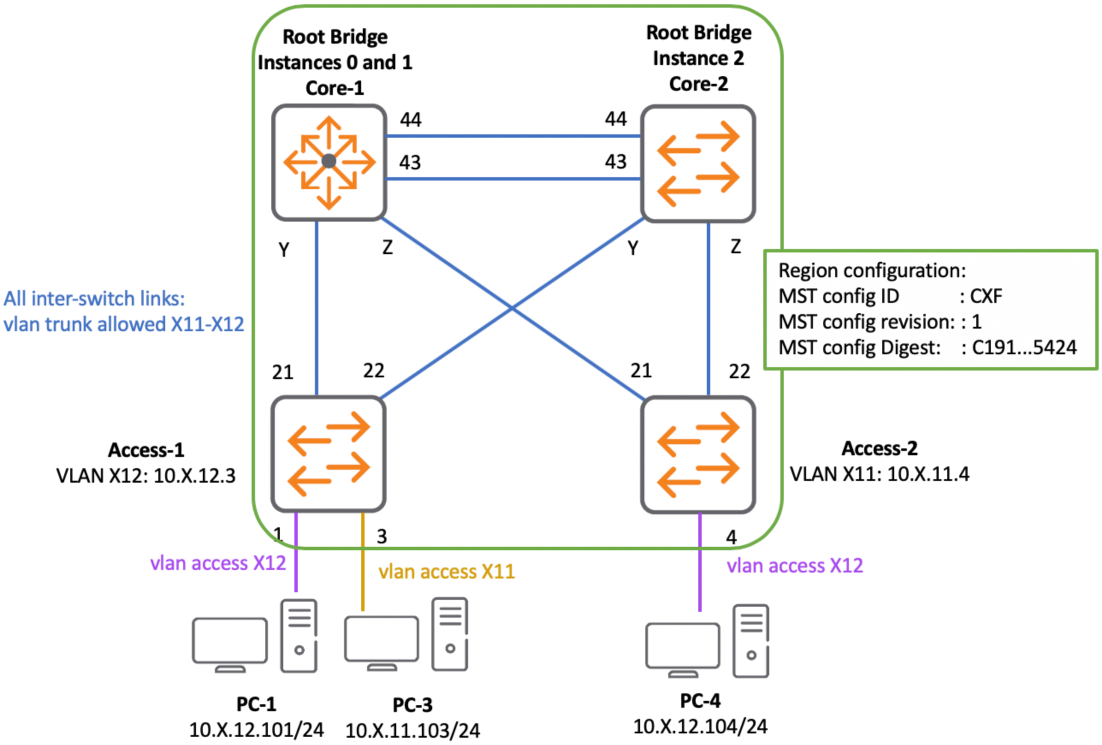

Figure 5.4-1: Lab Topology

Task 1: Port Priority Manipulation

Objectives

After deploying the MST region configuration properly on all switches, load sharing starts working across both uplinks. However, the link between the Core switches on port 44 is always inactive. This is because between ports 43 and 44, the latest has the highest Port ID value in all three instances.

In Instance 1 port 43 is Root and 44 is Alternate on Core-2. In Instance 2 port 43 is Root and 44 is Alternate on Core-1’s side. In this activity the instructor (Rent4Cheap Properties network administrator) will manipulate the port ID on one of them in order to invert the port’s roles in Instances 0 and 1.

> **NOTE:** Although this is an instructor-based demonstration, step-by-step instructions are provided in case you want to try this out in a home lab.

Steps

Core-2 (via PC-1)

1. Move to **Core-2**.
2. Use the “**show spanning-tree mst 1**” command and look for interfaces 1/1/43 and 1/1/44.

```
Core-2(config)# show spanning-tree mst 1 | begin 32 Role
Port           Role           State        Cost       Priority   Type
-------------- -------------- ------------ ---------- ---------- ----------
                            ←---- output omitted ---→
1/1/41         Disabled       Blocking     20000      128        point_to_point
1/1/43         Root           Forwarding   20000      128        point_to_point
1/1/44         Alternate      Blocking     20000      128        point_to_point
1/1/46         Disabled       Blocking     20000      128        point_to_point
                            ←---- output omitted ---→
Core-2(config)#
```

3. Repeat **step 2** for instance 2.

```
Core-2(config)# show spanning-tree mst 2 | begin 32 Role
Port           Role          State        Cost       Priority   Type
-------------- ------------- ------------ ---------- ---------- ----------
                             ←---- output omitted ---→
1/1/41         Disabled      Blocking     20000      128        point_to_point
1/1/43         Designated    Forwarding   20000      128        point_to_point
1/1/44         Designated    Forwarding   20000      128        point_to_point
1/1/46         Disabled      Blocking     20000      128        point_to_point
                             ←---- output omitted ---→
Core-2(config)#
```

What are the roles of port 1/1/43 on instances 1 and 2?

Instance 1: ____

Instance 2: ____

What are the roles of port 1/1/44 on instances 1 and 2?

Instance 1: ____

Instance 2: ____

TIP: Since Core-1 is Root on instance 0 and 1, you can assume that the port roles on instance 0 and 1 are the same. There is no need for verification.

IMPORTANT: Port priority is a value that in combination with the port number results in the Port ID. Port priority comes in a range between 0 to 240 and is configurable in AOS-CX switches in steps (multiplying factors) of 16 where the default value is step 8 (128) e.g. 128:43 and 128:4 in the output above.

If you remember from **Lab 5.1 Task 2 Step 13** Rule 8 stays: "*If a non-designated bridge has two or more links with equal RPC to the same Designated Bridge, then the local interface that connects neighbor's with lowest Port ID will be selected Root Port.*"

Each Designated bridge communicates its outbound interface's Port ID by including the information within the STP BPDU it transmits or relays to the non-Designated Bridge on that particular interface.

Here are two examples of how port priority and port ID work:

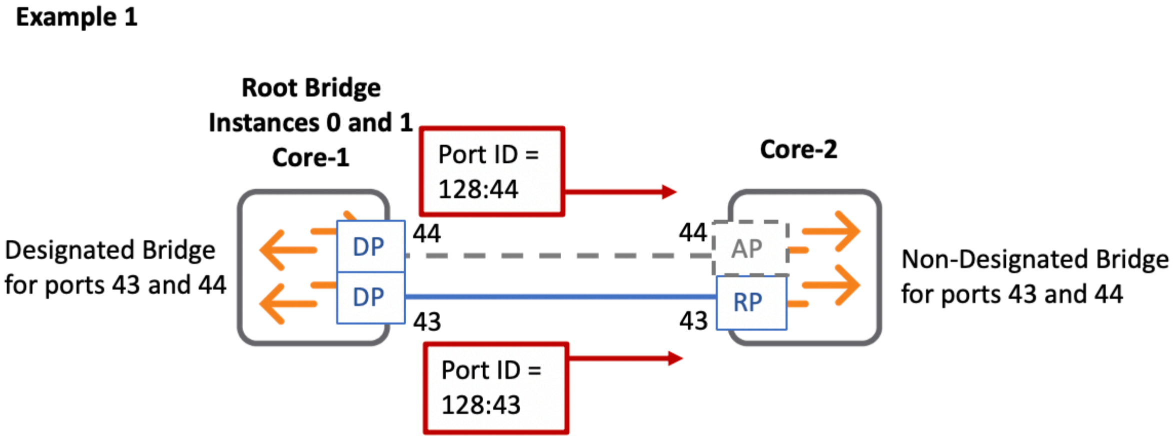

In first example Core-1 is the Root Bridge for Instance 0 and 1. It is sending BPDUs on ports 43 and 44 with Port IDs 128:43 and 128:44 respectively. When these messages arrive at Core-2, Core-2 chooses the local port 43 as its root port because it is connected to the interface with lowest Port ID on Core-1 side while local port 44 is set as alternate.

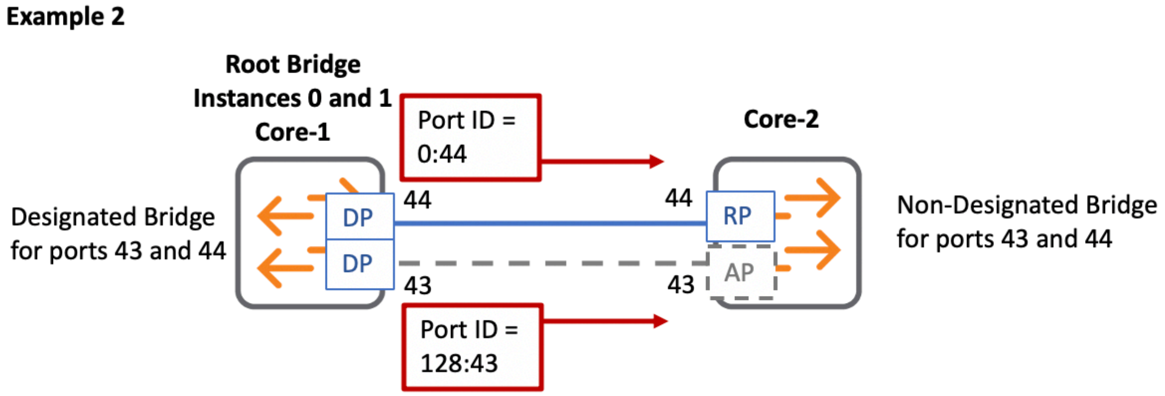

On this second example port priority has been changed on port 44 in Core-1. It now sends BPDUs with Port IDs 128:43 and 0:44. When these messages arrive at Core-2, Core-2 chooses local port 44 as Root port because it is connected to the interface with lowest Port ID on the Core-1 side while local port 43 is set as alternate.

Now proceed and deploy the solution from example 2.

Core-1 (via PC-1)

4. Move to **Core-1**.
5. Use the "**show spanning-tree mst 1**" command and look for interfaces 1/1/43 and 1/1/44.

```
Core-1(config)# show spanning-tree mst 1 | begin 32 Role
Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
←---- output omitted ---→
1/1/41         Disabled       Blocking     20000   128        point_to_point
1/1/43         Designated     Forwarding   20000   128        point_to_point
1/1/44         Designated     Forwarding   20000   128        point_to_point
1/1/46         Disabled       Blocking     20000   128        point_to_point
←---- output omitted ---→
Core-1(config)#
```

6. Repeat **step 5** for **instance 2**.

```
Core-1(config)# show spanning-tree mst 2 | begin 32 Role
Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
←---- output omitted ---→
1/1/41         Disabled       Blocking     20000   128        point_to_point
1/1/43         Root           Forwarding   20000   128        point_to_point
1/1/44         Alternate      Blocking     20000   128        point_to_point
1/1/46         Disabled       Blocking     20000   128        point_to_point
←---- output omitted ---→
Core-1(config)#
```

What are the roles of port 1/1/43 on instances 1 and 2?

Instance 1: ____

Instance 2: ____

What are the roles of port 1/1/44 on instances 1 and 2?

Instance 1: ____

Instance 2: ____

What is the port priority of ports 1/1/43 and 1/1/44 on instances 1 and 2?

____

NOTE: Notice the role port 1/1/43 has on Core-2 in instance 1, is the same role Core-1 has for that same port in Instance 2 (Root), likewise the role port 1/1/44 has on Core-2 Instance 1 is the same role Core-1 has in Instance 2 (Alternate). This is about to change.

7. Change **port 1/1/44**'s priority to 0 in **Instances 0 and 1**.

```
Core-1(config)# interface 1/1/44
Core-1(config-if)# spanning-tree port-priority 0
Core-1(config-if)# spanning-tree instance 1 port-priority 0
Core-1(config-if)#
```

Core-2 (via PC-1)

8. Move to **Core-2**.
9. Use the "**show spanning-tree mst 1**" command and look for interfaces **1/1/43 and 1/1/44**.

```
Core-2(config)# show spanning-tree mst 1 | begin 32 Role
Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
1/1/41         Disabled       Blocking     20000   128        point_to_point
1/1/43         Alternate      Blocking     20000   128        point_to_point
1/1/44         Root           Forwarding   20000   128        point_to_point
1/1/46         Disabled       Blocking     20000   128        point_to_point
Core-2(config)#
```

What are the roles of port 1/1/43 on instances 1 and 2?

Instance 1: ____

Instance 2: ____

Are both links active for at least one instance?

____

NOTE: The MST topology has now changed for Instance 0 and 1. It looks like figure 5.4-2, while topology for Instance 2 remains as before (figure 5.4-3).

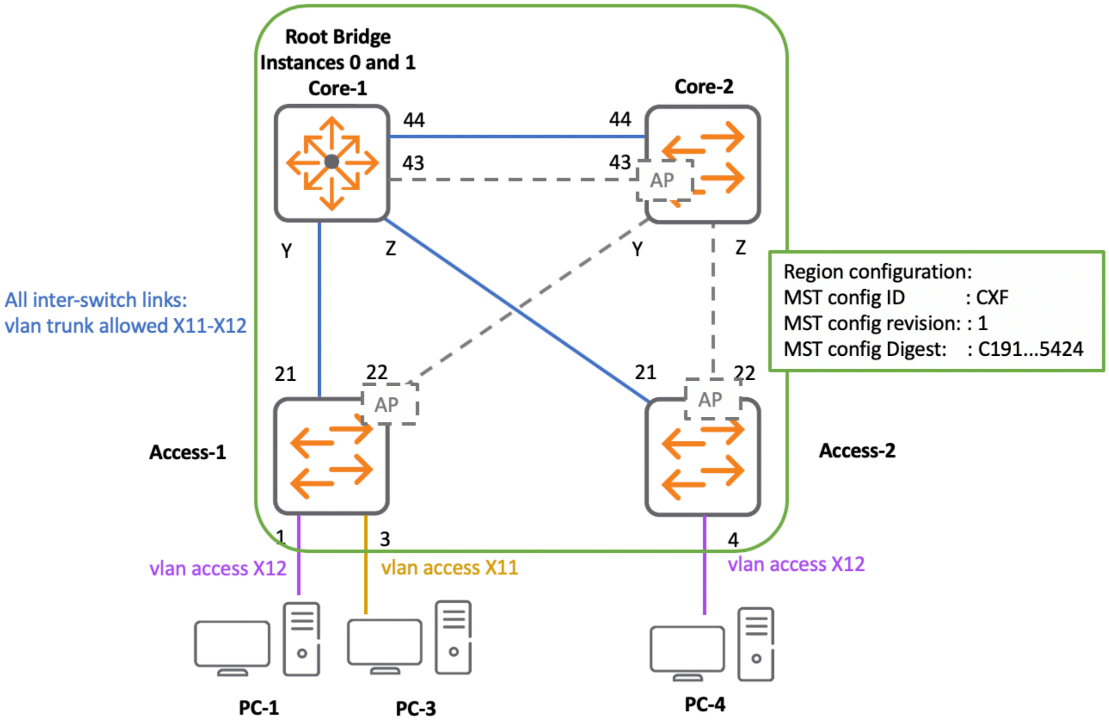

Figure 5.4-2: New topology of instances 0 and 1.

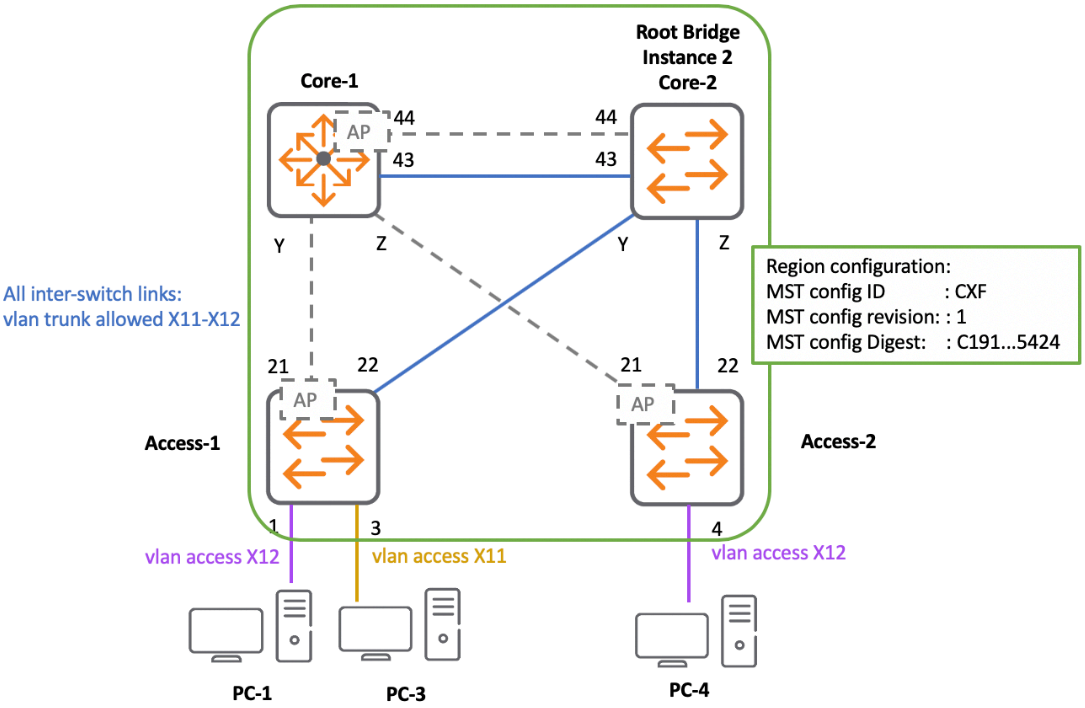

Figure 5.4-3: Topology of instance 2.

Task 2: Save Your Configurations

Objectives

Save your configurations and create checkpoints. Note that lab checkpoints might be used by later activities.

Steps

Core-1 and Core-2 (via PC-1).

1. Save the current Access and Core switches' configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

2. End the SSH sessions.

```
Core-1# exit
```

```
Core-2# exit
```

You have completed Lab 5.4!

AOS-CX Switching Fundamentals

Lab 6.1: Link Aggregation between Core Switches

Overview

After successfully deploying MST based load sharing on links between Core switches, the network administrator of Rent4Cheap Properties has been monitoring the bandwidth utilization of links of ports 43 and 44 and calculated an average of 10% utilization of one link versus 55% in the other. Although neither link is congested yet, the network administrator would like to look for a better way share the load among links.

Although moving VLANs from one instance to the other looks like a good solution and might work in the short term, this is not a scalable option because nothing guarantees that traffic patterns will not change tomorrow, in a week, or a few months from now.

The network administrator has approached you and asked for advice. You propose deploying link aggregation, since load sharing is not VLAN based but hash based (based on layer 2 or layer 3 source and destination addresses) which commonly leads to more even resource utilization.

NOTE: Although this is an instructor-based demonstration, the steps are provided in case you want to try this out yourself in a private lab.

Objectives

After completing this lab, you will be able to:

- Deploy static Link Aggregation
- Understand the nature of transient loops when creating static aggregations
- Monitor LAG interfaces in AOS-CX

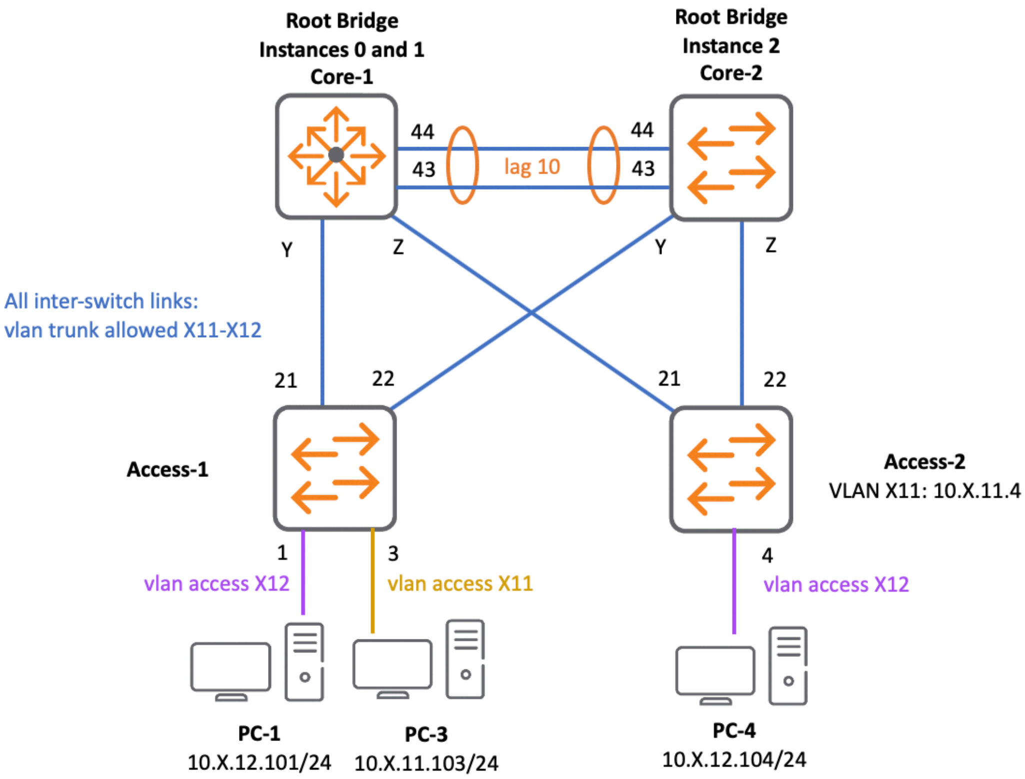

Figure 6.1-1: Lab Topology

Task 1: Pre-lab Setup:

Objectives

In this activity you will load Lab5-2_final checkpoint in Access-1 and Access-2, where those two switches were interconnected to the Core switches using ports 1/1/21 and 1/1/22.

> **NOTE:** This activity is dependent on Lab 5.2 configuration, make sure you have completed that lab before starting the current one. **Do not proceed if this is not the case**.

Steps

Access-1 and Access-2

1. Display the checkpoint list and confirm the Lab5-2_final checkpoint is there.

```
T11-Access-1# show checkpoint list | include Lab5
Lab5-1_final
Lab5-2_final
Lab5-3_final
Lab5-3_task1_done
T11-Access-1#
```

```
T11-Access-2# show checkpoint list | include Lab5
Lab5-1_final
Lab5-2_final
Lab5-3_final
Lab5-3_task1_done
T11-Access-2#
```

2. Load the checkpoint using the **checkpoint rollback** command.

```
T11-Access-1# checkpoint rollback Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# checkpoint rollback Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

Task 2: Configure Manual Link Aggregation

Objectives

The network administrator of Rent4Cheap Properties (your instructor) will demonstrate and test out static aggregation on the links between the Core switches. He researched the configuration commands and is ready to add them during a maintenance window.

Steps

PC-3

1. Access PC-3.
2. Run a continuous ping to the IP address of Access-2 on **VLAN X11** (**10.X.11.4**). The ping should be successful.

Core-1 (via PC-1)

3. Open the SSH session to Core-1.
4. Create LAG 10 interface and apply a description. This will be used as a logical Layer 2 connection between Cores.

```
Core-1# configure terminal
Core-1(config)# interface lag 10
Core-1(config-lag-if)# description TO_CORE-2_PORTS-43_&_44
Core-1(config-lag-if)#
```

5. Disable routing and enable the interface.

```
Core-1(config-lag-if)# no routing
Core-1(config-lag-if)# no shutdown
Core-1(config-lag-if)#
```

6. Allow VLANs **X11** and **X12**.

```
Core-1(config-lag-if)# vlan trunk allowed X11-X12
Core-1(config-lag-if)#
```

7. Create a port range with interfaces **1/1/43** and **1/1/44** and make these two ports members of **LAG 10**.

```
Core-1(config-if)# interface 1/1/43-1/1/44
Core-1(config-if-<1/1/43-1/1/44>)# lag 10
Core-1(config-if-<1/1/43-1/1/44>)# end
```

8. Display detailed information about **LAG 10**.

```
Core-1# show interface lag 10


 Aggregate lag10 is up
 Admin state is up
 Description : TO_CORE-2_PORTS-43_&_44
 MAC Address                  : 90:20:c2:bc:ed:00
 Aggregated-interfaces        : 1/1/43 1/1/44
 Aggregation-key              : 10
 Speed                        : 2000 Mb/s
 L3 Counters: Rx Disabled, Tx Disabled
 qos trust none
 VLAN Mode: native-untagged
 Native VLAN: 1
 Allowed VLAN List: 1111-1112
 Rx
       164524 input packets       24564243 bytes
            0 input error             2766 dropped
            0 CRC/FCS
 Tx
       276626 output packets      46476581 bytes
            0 input error             2774 dropped
            0 collision
Core-1#
```

What is the state of LAG 10?

What are the member ports?

____

What is the speed of the link?

____

How is that speed determined?

____

What VLANs are forwarding traffic on this lag?

____

How many packets are being sent and received?

____

Are all these packets generated by the continuous ping you are running?

____

NOTE: Right now, interface LAG 10 is up because the previous configuration has created a local static aggregation that does not depend on any control plane protocol-based negotiation with the remote end (Core-2). However, this has data plane implications, the number of sent and received packets are not the result of a continuous ping. The question is: what else can be creating that amount of traffic? After all, you are in the middle of a maintenance window and nobody else is working in the network.

PC-3

9. Move back to PC-3.

Is the ping still running?

____

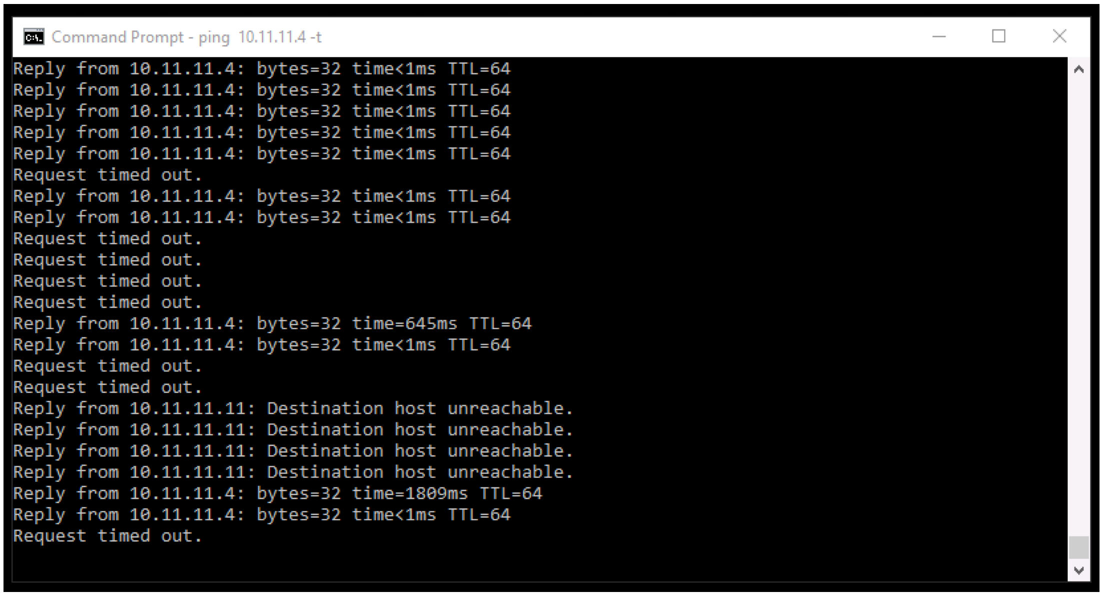

```
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Request timed out.
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.11.4: bytes=32 time=645ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.4: bytes=32 time=1809ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Request timed out.
```

Figure 6.1-2: Failed ping

NOTE: Please be patient on getting figure 6.1-2 result. Since the lab environment has just a few stations, and not all of them are sending broadcast traffic, it may take some time before pings fail. In fact it is possible that ICMP traffic does not fail at all, especially if Core switches are 8300s. Nonetheless this issue shows up right away on a production network, where hundreds of endpoints are connected.

NOTICE: You are experiencing a transient layer 2 loop. When you configured static link aggregation, on Core-1 it started sending every single frame to Core-2 on either port 43 or 44 based on a load sharing mechanism that uses Source and Destination IPs (or Source and Destination MACs in absence of IP headers) as input and gives a hash result as output: either 0 or 1 that represent port 43 and 44 respectively. This includes the BPDUs, since at the STP level, LAG 10 is a single logical port.

Core-2 is not running static aggregation yet, and its STP processes see two physical ports instead of one and, Core-2 only receives BPDUs on one of these ports. After a few seconds, the lack of BPDUs in one port forces it to transition its role to Designated (as if it was an interface connected to an endpoint) while the other interface becomes Root, these events happen on Instance 0 and 1, because on instance 2 both ports on Core-2 are already Designated.

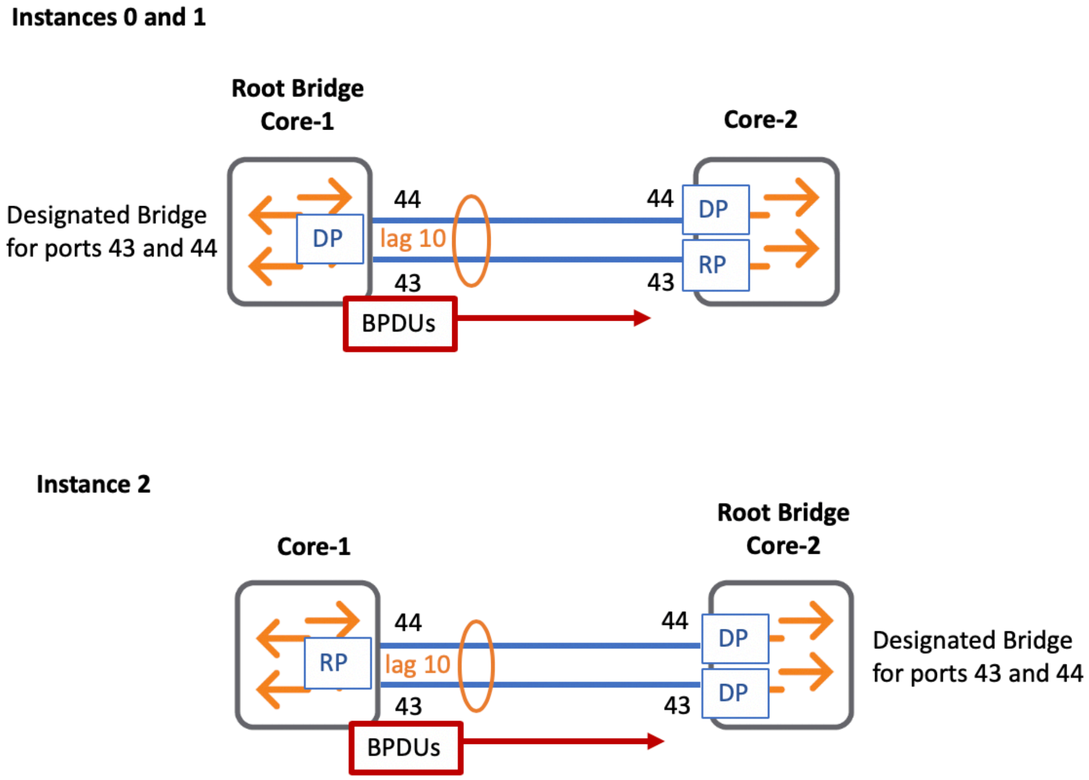

Figure 6.1-3: BPDUs in static aggregation

In both cases Core-2's ports eventually move to Forwarding mode. The problem appears when Core-1 forwards a broadcast, multicast or unknown unicast frame across the lag. It uses one of the physical links, and when Core-2 receives it, it forwards the traffic to all interfaces the VLAN belongs to, including the second link back to Core-1.

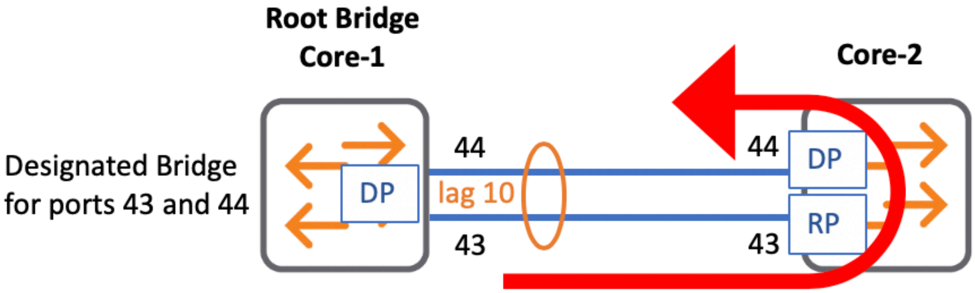

Figure 6.1-4: Transient loop

This means Core-2 is looping some frames back to Core-1, including the BPDUs it relays. This causes MAC address flapping. Also, each broadcast or multicast Core-2 receives in its downlinks will be sent across both port 43 and 44 generating duplicated frames. That is why pings either stopped or were inconsistent.

The solution is to disable one of the ports (preferably the former Alternate port), before starting the static aggregation configuration, and re-enable it once it is done on both sides.

IMPORTANT: Another potential loop situation can take place when configuring Static Aggregation in Access switches uplinks that terminate on different non-related/non-stacked physical devices.

Therefore, before configuring static aggregation, you must verify the following:

1. All LAG member ports except one are disabled on one side.
2. Confirm cabling is correct and involves two switching entities only.

Since you are already facing the issue, you will begin by removing the transient loop, then you will complete Core-2's portion of the setup.

Core-2 (via PC-1)

10. Move to Core-2.

11. Disable port 1/1/43.

```
Core-2# configure terminal
Core-2(config)# interface 1/1/43
Core-2(config-if)# shutdown
Core-2(config-if)# exit
```

PC-3

12. Move back to PC-3.

Is the ping still failing?

___

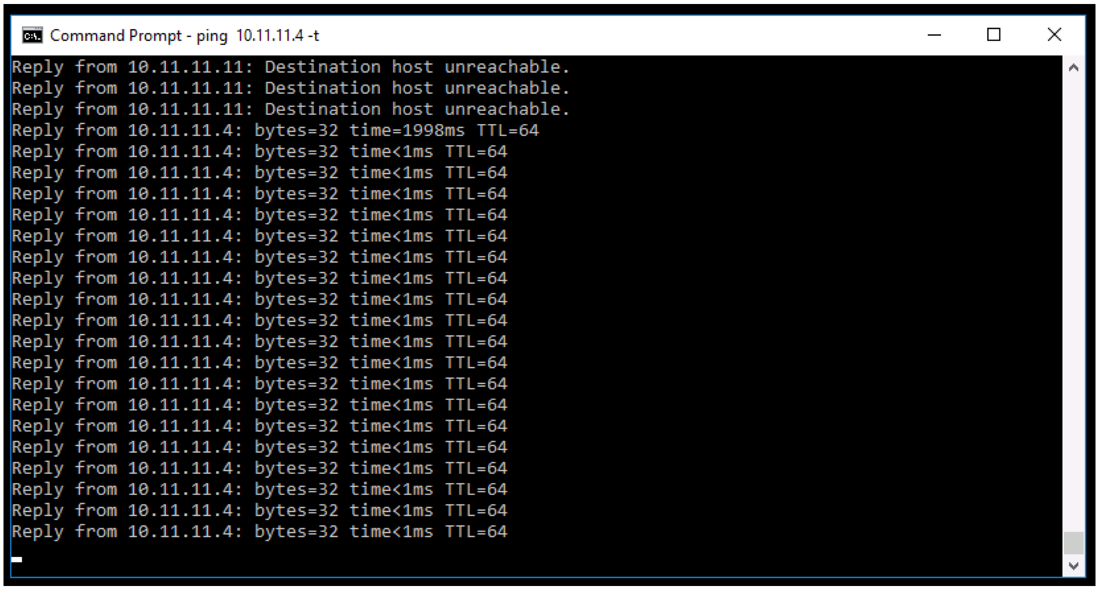

```
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.11: Destination host unreachable.
Reply from 10.11.11.4: bytes=32 time=1998ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
Reply from 10.11.11.4: bytes=32 time<1ms TTL=64
```

Figure 6.1-5: Ping running

Core-2

13. Move back to Core-2.
14. Repeat steps 4 to 8.

```
Core-2(config)# interface lag 10
Core-2(config-lag-if)# description TO_CORE-1_PORTS-43_&_44
Core-2(config-lag-if)# no routing
Core-2(config-lag-if)# no shutdown
```

```
Core-2(config-lag-if)# vlan trunk allowed X11-X12
Core-2(config-lag-if)#
```

```
Core-2(config-if)# interface 1/1/43-1/1/44
Core-2(config-if-<1/1/43-1/1/44>)# lag 10
Core-2(config-if-<1/1/43-1/1/44>)# exit
```

15. Enable interface **1/1/43** back.

```
Core-2(config)# interface 1/1/43
Core-2(config-if)# no shutdown
Core-2(config-if)# end
```

16. Show interface **LAG 10** status.

```
Core-2# show interface lag brief
--------------------------------------------------------------------------------
----
Port      Native  Mode   Type            Enabled Status  Reason
Speed
          VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
lag10     1       trunk  --              yes     up      --
2000
Core-2#
```

Is LAG 10 working normally?

____

PC-3

17. Move back to PC-3.

Is the ping still working?

____

Task 3: Normalize Configuration for All Kits.

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

PC-1

1. Access PC-1
2. Open the "**CXF LAG 10.txt**" file located within CXF Student Folder on the desktop.

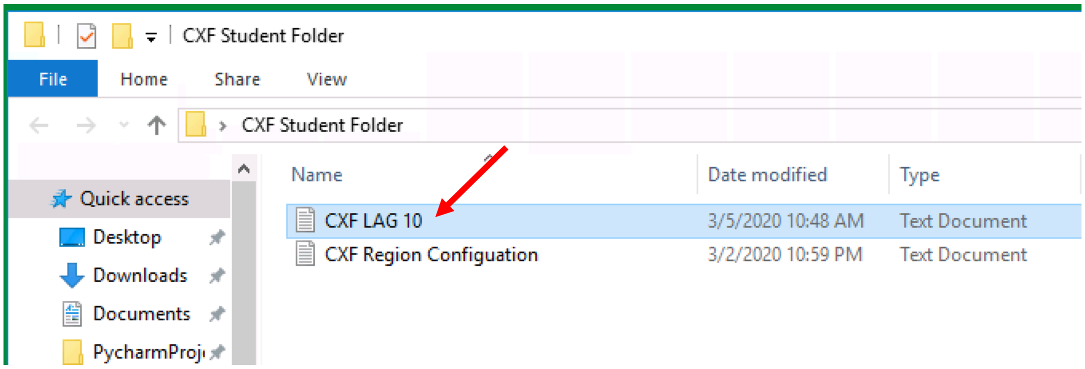

Name	Date modified	Type
CXF LAG 10	3/5/2020 10:48 AM	Text Document
CXF Region Configuation	3/2/2020 10:59 PM	Text Document

3. Copy the file contents into the clipboard.

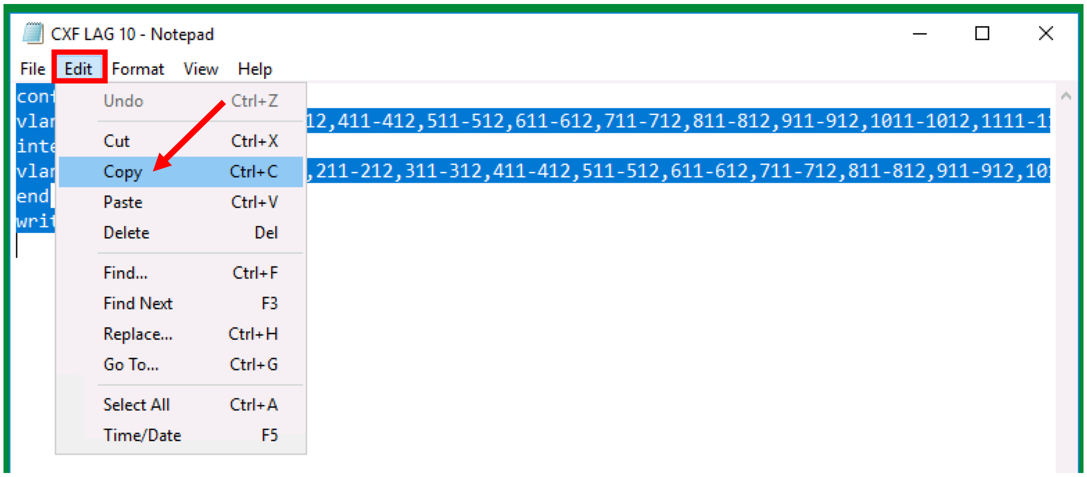

Core-1 and Core-2 (via PC-1)

4. Paste contents on Core-1 and Core-2's SSH sessions.

```
Core-1# end
Core-1# configure terminal
Core-1(config)# vlan 111-112,211-212,311-312,411-412,511-512,611-612,711-712,811-
812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-1412
Core-1(config-vlan-<111-112,211-212,311-312,411-412,511-512,611-612,7>)#
interface lag 10
Core-1(config-lag-if)# vlan trunk allowed 111-112,211-212,311-312,411-412,511-
512,611-612,711-712,811-812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-
1412
Core-1(config-lag-if)# end
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# end
Core-2# configure terminal
Core-2(config)# vlan 111-112,211-212,311-312,411-412,511-512,611-612,711-712,811-
812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-1412
Core-2(config-vlan-<111-112,211-212,311-312,411-412,511-512,611-612,7>)#
interface lag 10
Core-2(config-lag-if)# vlan trunk allowed 111-112,211-212,311-312,411-412,511-
512,611-612,711-712,811-812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-
1412
Core-2(config-lag-if)# end
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

NOTE: You may notice that the console got overrun due the length of some commands in the script, nonetheless it should not prevent the system from applying the configuration and you should not get a warning.

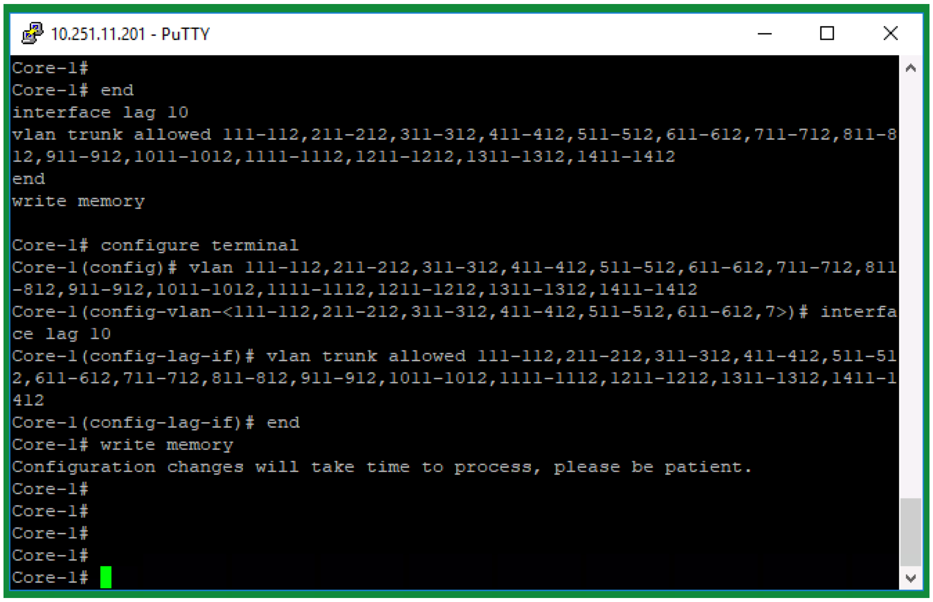

```
Core-1#
Core-1# end
interface lag 10
vlan trunk allowed 111-112,211-212,311-312,411-412,511-512,611-612,711-712,811-8
12,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-1412
end
write memory

Core-1# configure terminal
Core-1(config)# vlan 111-112,211-212,311-312,411-412,511-512,611-612,711-712,811
-812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-1412
Core-1(config-vlan-<111-112,211-212,311-312,411-412,511-512,611-612,7>)# interfa
ce lag 10
Core-1(config-lag-if)# vlan trunk allowed 111-112,211-212,311-312,411-412,511-51
2,611-612,711-712,811-812,911-912,1011-1012,1111-1112,1211-1212,1311-1312,1411-1
412
Core-1(config-lag-if)# end
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
Core-1#
Core-1#
Core-1#
Core-1#
```

5. End the SSH sessions.

```
Core-1# exit
```

```
Core-2# exit
```

You have completed Lab 6.1!

AOS-CX Switching Fundamentals

Lab 6.2 Deploying LACP based Link Aggregation

Overview

When LAG 10 was created between both Core switches, BigStartup saw the value of the technology and asked about other potential use cases. When you mentioned link aggregations can be used between switches, routers, firewalls, and servers, the customer became more interested. They asked if it is possible to deploy aggregated links without any chance of loops and can you demonstrate the technology?

Objectives

After completing this lab, you will be able to:

- Deploy LACP based Link Aggregation
- Demonstrate the benefits of LACP vs Static aggregation

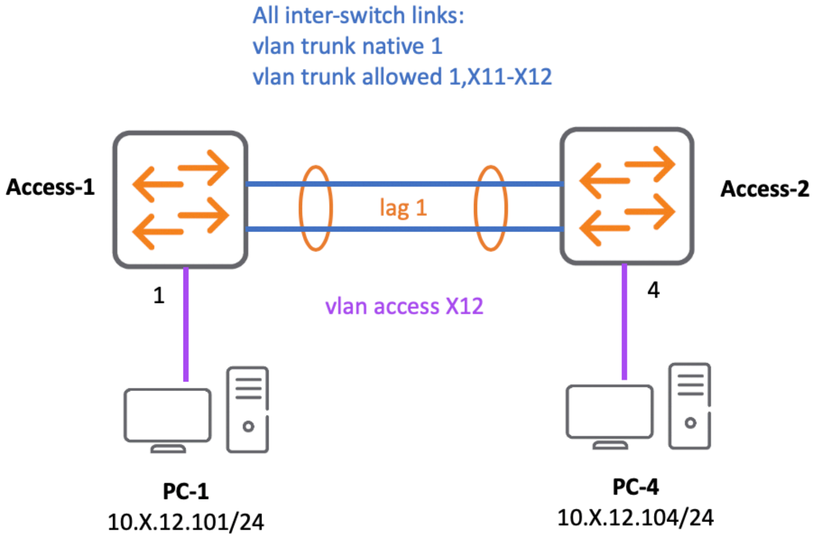

Figure 6.2-1: Lab Topology

Task 1: Pre-lab Setup

Objectives

In this activity you will isolate Access-1 and Access-2 from the rest of the network then enable a dual homed topology using ports 27 and 28 and create an aggregated link LAG 1 using those two ports.

NOTE: This activity uses the same topology you created in Lab 5.3 Task 1 configuration, if you completed that task then feel free to load “Lab5-3_task1_done” checkpoints on both Access switches and move to Task 2.

```
T11-Access-1# checkpoint rollback Lab5-3_task1_done
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# checkpoint rollback Lab5-3_task1_done
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

If this is not the case, then complete steps 1 to 7.

Steps

Access-1

1. Open a console connection to Access-1. Login using **admin** and **no password**.
2. Disable ports 1/1/21 and 1/1/22.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface 1/1/21-1/1/22
T11-Access-1(config-if-<1/1/21-1/1/22>)# shutdown
T11-Access-1(config-if-<1/1/21-1/1/22>)# exit
```

3. Access port **1/1/27** and set a description.

```
T11-Access-1(config)# interface 1/1/27
T11-Access-1(config-if)# description TO_TX-ACCESS-2_PORT-27
T11-Access-1(config-if)# exit
```

4. Create a port range including 1/1/27 and 1/1/28, allow VLANs 1, X11 and X12, then enable them.

```
T11-Access-1(config)# interface 1/1/27-1/1/28
T11-Access-1(config-if-<1/1/27-1/1/28>)# vlan trunk allow 1,X11-X12
T11-Access-1(config-if-<1/1/27-1/1/28>)# no shutdown
T11-Access-1(config-if-<1/1/27-1/1/28>)# exit
```

5. Confirm port 1/1/21 to 1/1/22 are down.

```
T11-Access-1(config) # show interface 1/1/21-1/1/22 link-status
--------------------------------------------------
Port      Type            Physical    Link
                          Link State  Transitions
--------------------------------------------------
1/1/21    1GbT            down        4
1/1/22    1GbT            down        2
T11-Access-1(config) # exit
```

NOTICE: Remember that you are about to create a Layer 2 loop, which has the potential of affecting other students. In order to limit the effects, you have to make sure that both uplinks 1/1/21 and 1/1/22 are down. **If this is not the case, go to those ports and shut them down.**

Access-2

6. Move to Access-2, then repeat steps 2 through 4.

```
T11-Access-2# configure terminal
T11-Access-2(config-if-<1/1/21-1/1/22>)# interface 1/1/21-1/1/22
T11-Access-2(config-if-<1/1/21-1/1/22>)# shutdown
T11-Access-2(config-if-<1/1/21-1/1/22>)# exit
```

```
T11-Access-2(config)# interface 1/1/27
T11-Access-2(config-if)# description TO_TX-ACCESS-1_PORT-27
T11-Access-2(config-if)# exit
```

```
T11-Access-2(config)# interface 1/1/27-1/1/28
T11-Access-2(config-if-<1/1/27-1/1/28>)# vlan trunk allow 1,X11-X12
T11-Access-2(config-if-<1/1/27-1/1/28>)# no shutdown
T11-Access-2(config-if-<1/1/27-1/1/28>)# exit
```

7. Confirm ports 1/1/21 to 1/1/22 are down and 1/1/27 to 1/1/28 are up.

```
T11-Access-2(config)# show interface 1/1/21-1/1/22,1/1/27-1/1/28 link-status
--------------------------------------------------
Port      Type            Physical    Link
                          Link State  Transitions
--------------------------------------------------
1/1/21    1GbT            down        0
1/1/22    1GbT            down        4
1/1/27    SFP+DAC1        up          11
1/1/28    SFP+DAC1        up          17
T11-Access-2(config)#
```

NOTICE: Remember that you are about to create a Layer 2 loop, which has the potential of affecting the entire network, in order to limit the effects, you have to make sure that both uplinks 1/1/21 and 1/1/22 are down. **Do not proceed if this is not the case.**

Task 2: Configure LACP Link Aggregation

Objectives

In the current task you will deploy an aggregated link between both Access Switches using LACP for negotiating the physical ports' states.

Steps

Access-1

1. Open a console connection to Access-1.
2. Create LAG 1 and add a description.

```
T11-Access-1# configure terminal
T11-Access-1(config)# interface lag 1
T11-Access-1(config-lag-if)# description LAG_TO_Access-2
```

3. Run Active LACP and fast rate heartbeats on the link aggregation.

```
T11-Access-1(config-lag-if)# lacp mode active
T11-Access-1(config-lag-if)# lacp rate fast
T11-Access-1(config-lag-if)#
```

4. Allow VLANs 1, X11 and X12. then enable the interface.

```
T11-Access-1(config-lag-if)# vlan trunk allowed 1,X11,X12
T11-Access-1(config-lag-if)# no shutdown
T11-Access-1(config-lag-if)# exit
```

5. Make ports 1/1/27 and 1/1/28 part of the lag.

```
T11-Access-1(config)# interface 1/1/27-1/1/28
T11-Access-1(config-if-<1/1/27-1/1/28>)# lag 1
T11-Access-1(config-if-<1/1/27-1/1/28>)# exit
```

6. Use the **show lacp configuration** command for displaying the local System-ID and Priority.

```
T11-Access-1(config)# show lacp configuration

System-ID       : 88:3a:30:98:30:00
System-priority : 65534
T11-Access-1(config)#
```

7. Display the LACP based LAG status information.

```
T11-Access-1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf    Aggr       Port  Port  State   System-ID         System Aggr Forwarding
        Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/27  lag1       28    1     ASFOE   88:3a:30:98:30:00 65534  1    lacp-block
1/1/28  lag1       29    1     ASFOE   88:3a:30:98:30:00 65534  1    lacp-block


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf    Aggr       Port  Port  State   System-ID         System Aggr
        Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/27  lag1       0     65534 PLFOEX  00:00:00:00:00:00 65534  0
1/1/28  lag1       0     65534 PLFOEX  00:00:00:00:00:00 65534  0
T11-Access-1(config)#
```

What is the local system ID?

____

What is the remote system ID?

What is the forwarding state?

ANSWER: Forwarding state is lacp-block this prevents data packets from being transmitted on such physical ports until the local switch receives inbound LACP Data Units from a peer, preventing any transient loops.

Access-2

8. Open a console connection to Access-2.
9. Repeat steps 2 to 5 using VLANs X11 and X12.

```
T11-Access-2# configure terminal
T11-Access-2(config)# interface lag 1
T11-Access-2(config-lag-if)# description LAG_TO_Access-1
```

```
T11-Access-2(config-lag-if)# lacp mode active
T11-Access-2(config-lag-if)# lacp rate fast
T11-Access-2(config-lag-if)#
```

```
T11-Access-2(config-lag-if)# vlan trunk allowed 1,X11,X12
T11-Access-2(config-lag-if)# no shutdown
T11-Access-2(config-lag-if)# exit
```

```
T11-Access-2(config)# interface 1/1/27-1/1/28
T11-Access-2(config-if-<1/1/27-1/1/28>)# lag 1
T11-Access-2(config-if-<1/1/27-1/1/28>)# exit
```

10. Display the LACP based LAGs status information.

```
T11-Access-2(config)# show lacp interfaces
```

```
State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf     Aggr       Port  Port  State   System-ID         System Aggr Forwarding
         Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/27   lag1       28    1     ASFNCD  88:3a:30:97:a4:40 65534  1    up
1/1/28   lag1       29    1     ASFNCD  88:3a:30:97:a4:40 65534  1    up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf     Aggr       Port  Port  State   System-ID         System Aggr
         Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/27   lag1       28    1     ASFNCD  88:3a:30:98:30:00 65534  1
1/1/28   lag1       29    1     ASFNCD  88:3a:30:98:30:00 65534  1
T11-Access-2(config)#
```

What physical ports are a member of the LAG?

____

What are the state flags on the local and remote ports?

____

____

What is their meaning?

____

____

____

____

11. Issue the **show spanning-tree** command.

```
T11-Access-2(config)# show spanning-tree | exclude Disabled

                              ←---- output omitted ---→

Port          Role           State        Cost           Priority   Type
------------ -------------- ------------ -------------- ---------- ----------
1/1/4         Designated     Forwarding   20000          128        point_to_point
lag1          Root           Forwarding   2000           64         point_to_point

T11-Access-2(config)#
```

What is the spanning tree state of ports 1/1/27, 1/1/28 and LAG 1?

____

ANSWER: Ports 1/1/27 and 1/1/28 are not listed, while lag1 is Root. When LAG 1 was created and ports 1/1/27 and 1/1/28 became part of it, then Spanning Tree stopped considering the physical interfaces in its calculations and started using LAG 1 instead.

12. Run the **show lacp aggregates** command.

```
T11-Access-2(config)# show lacp aggregates

Aggregate name    : lag1
Interfaces        : 1/1/28 1/1/27
Heartbeat rate    : Fast
Hash              : l3-src-dst
Aggregate mode    : Active

T11-Access-2(config)#
```

What is the current (default) hashing algorithm?

____

PC-1

13. Open a console session to PC-1.
14. Ping PC-4 (10.X.12.104). Ping should be successful.

NOTE: Since this traffic will always have the same source and destination IP addresses, only one link is being used for sending the traffic in either direction. If multiple clients were connected on both switches, then the traffic between them would be shared across both links in an Active/Active way.

Task 3: Save Your Configurations

Objectives

You will now proceed to create checkpoints.

Steps

Access-1 and Access-2

1. Save the current Access switches' configuration in the startup checkpoint.

```
T11-Access-1# write memory
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# write memory
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

2. Backup the current Access switches' configuration as a custom checkpoint called **Lab6-2_final**.

```
T11-Access-1# copy running-config checkpoint Lab6-2_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# copy running-config checkpoint Lab6-2_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

You have completed Lab 6.2!

AOS-CX Switching Fundamentals

Lab 7.1: IPv4 Routing

Overview

As the network grows, BigStartup has realized the need for communications between departments. Services such as Zoom conferencing, Remote Printing, Remote Assistance, and Internet access move traffic across VLANs. To provide for this new requirement, you have suggested enabling inter VLAN routing rather than reverting to a single VLAN design. This enables the connectivity level your customer is looking for and allows for blocking forbidden connection attempts using traffic filters (Routed Access Control Lists).

You will enable Layer 3 functions on one of your Core switches. Then the TCP IP stack on each client and host will require a default gateway IP address to enable using Layer 3 functions to deliver the packets destined to non-local segments.

Objectives

After completing this lab, you will be able to:

- Assign IP addresses to SVIs
- Enable Inter-VLAN routing
- Run traffic analysis using Wireshark
- Describe the end to end packet delivery

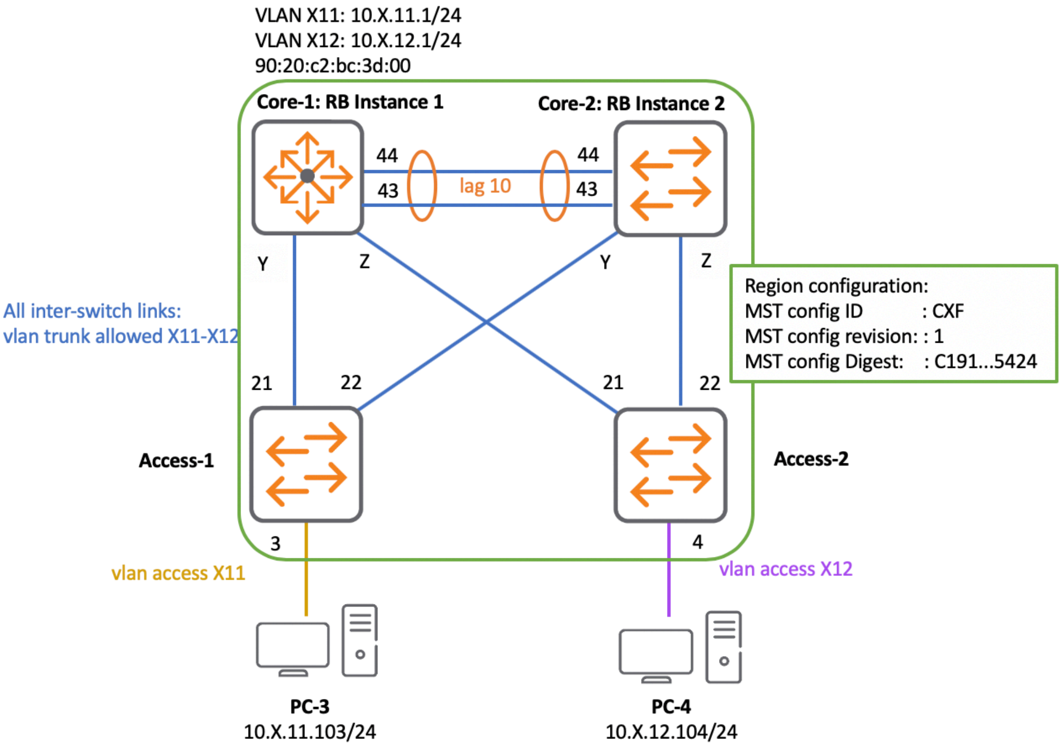

Figure 7.1-1: Lab Topology

Task 1: Pre-lab Setup:

Objectives

In this activity you will load Lab5-2_final checkpoint in Access-1 and Access-2, where those two switches were interconnected to the Core switches using ports 1/1/21 and 1/1/22.

> **NOTE:** This activity has a dependency on the Lab 5.2 configuration, make sure you have completed that lab before starting the current one. **Do not proceed if this is not the case**.

Steps

Access-1 and Access-2

1. Display the checkpoint list and confirm Lab5-2_final is there.

```
T11-Access-1# show checkpoint list | include Lab5
Lab5-1_final
Lab5-2_final
Lab5-3_final
Lab5-3_task1_done
T11-Access-1#
```

```
T11-Access-2# show checkpoint list | include Lab5
Lab5-1_final
Lab5-2_final
Lab5-3_final
Lab5-3_task1_done
T11-Access-2#
```

2. Load the checkpoint using the **checkpoint rollback** command.

```
T11-Access-1# checkpoint rollback Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-1#
```

```
T11-Access-2# checkpoint rollback Lab5-2_final
Configuration changes will take time to process, please be patient.
T11-Access-2#
```

Task 2: Set IP Default-Gateway

Objectives

In this first task you will configure IP addresses of both interface VLAN X11 and X12 in Core-1, then you will assign those addresses as default Gateways on PC-3 and PC-4.

Steps

Core-1 (via PC-1)

1. Open an SSH session to Core-1. Login using **cxfX/aruba123**

> **NOTE:** Replace the highlighted “X” for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

2. **Create interface vlan X11, then assign it IP address 10.X.11.1/24.**

> **NOTE:** Replace the highlighted “X” for your student table number.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11
Core-1(config-if-vlan)# ip address 10.X.11.1/24
Core-1(config-if-vlan)# exit
```

> **IMPORTANT:** This makes Core-1 a multilayer switch capable of routing traffic into the 10.X.11.0/24 segment.

3. See the newly created SVI details using “**show ip interface vlanX11**”

```
Core-1(config)# show ip interface vlanX11
```

```
Interface vlan1111 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:ed:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.11.11.1/24
 L3 Counters: Rx Disabled, Tx Disabled

Core-1(config)#
```

What are the SVI state and MAC address?

____

4. Repeat steps 2 and 3 for interface vlan X12, with IP address 10.X.12.1/24

```
Core-1(config)#
Core-1(config)# interface vlan X12
Core-1(config-if-vlan)# ip address 10.X.12.1/24
Core-1(config-if-vlan)# end
```

```
Core-1# show ip interface vlanX12

Interface vlan1112 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:ed:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.11.12.1/24
 L3 Counters: Rx Disabled, Tx Disabled

Core-1#
```

IMPORTANT: This command is case sensitive, make sure to type lowercase "vlan" (lowercase) immediately followed by the VLAN number, e.g. "**show ip interface vlan1111**" for table 11.

What is the SVI MAC address?

____

NOTE: Both SVIs use the same MAC address (the system one), this does not create any conflict because they are in two different broadcast domains.

5. Display the IPv4 routing table and look for your newly added prefixes.

```
Core-1# show ip route

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

                                ←---- output omitted ---→
10.9.11.0/24, vrf default
        via  vlan911,  [0/0],  connected
10.9.11.1/32, vrf default
        via  vlan911,  [0/0],  local
10.9.12.0/24, vrf default
        via  vlan912,  [0/0],  connected
10.9.12.1/32, vrf default
        via  vlan912,  [0/0],  local
10.10.11.0/24, vrf default
        via  vlan1011,  [0/0],  connected
10.10.11.1/32, vrf default
        via  vlan1011,  [0/0],  local
10.10.12.0/24, vrf default
        via  vlan1012,  [0/0],  connected
10.10.12.1/32, vrf default
        via  vlan1012,  [0/0],  local
10.11.11.0/24, vrf default
        via  vlan1111,  [0/0],  connected
10.11.11.1/32, vrf default
        via  vlan1111,  [0/0],  local
10.11.12.0/24, vrf default
        via  vlan1112,  [0/0],  connected
10.11.12.1/32, vrf default
        via  vlan1112,  [0/0],  local
                                ←---- output omitted ---→
```

NOTE: Since this device is a shared resource, the output that you get out of this command may contain either more or less entries.

TIP: When the routing table is that long, you can either use filtered versions of the command (e.g. “show ip route | begin 7 10.X.11.0”) or you can use a prefix specific command:

```
Core-1# show ip route 10.X.11.0/24
Displaying ipv4 routes selected for forwarding
'[x/y]' denotes [distance/metric]
10.11.11.0/24, vrf default
        via  vlan1111,  [0/0],  connected
Core-1#
```

NOTE: There are 4 prefixes published in the routing table after assigning the IP addresses. The ones with prefix length 32 are considered local and reference the IP addresses just configured in the SVIs.

The /24 prefixes are the connected subnets discovered from having an interface with an IP in those segments.

IP prefixes are expressed using the following format:

```
PREFIX/PREFIX_LENGHT, vrf VRF_NAME
    via OUTBOUND_INTERFACE, [DISTANCE/METRIC], ROUTING_PROCESS
```

Notice: they all contain vrf “default”. VRF stands for Virtual Routing and Forwarding, which is the control plane virtual routing table the system is using for moving traffic at Layer 3 in the data plane. AOS-CX has two built-in VRFs: mgmt for management traffic and default for data traffic.

IMPORTANT: AOS-CX switches can support several virtual routing table instances that are used for keeping IP Prefixes separated into different Layer 3 logical routing domains. Under normal circumstances, control plane prefixes from one VRF cannot be shared with other VRFs and data plane traffic contained in one VRF cannot be forwarded to interfaces belonging to another VRF (unless explicit prefix leaking is intentionally enabled).

This feature is ideal in multitenancy environments like Data Centers, Service Provider networks, and Network as a Service environments such as Rent4Cheap Properties.

Currently Core-1 is able to move traffic from either IP segment. You will add the client gateways. Non-local traffic will be delivered to the local gateway using Layer 2 and then forwarded to the non-local destinations using Layer 3.

PC-3

6. Access PC-3.
7. Assign 10.X.11.1 as the default gateway in Lab NIC.

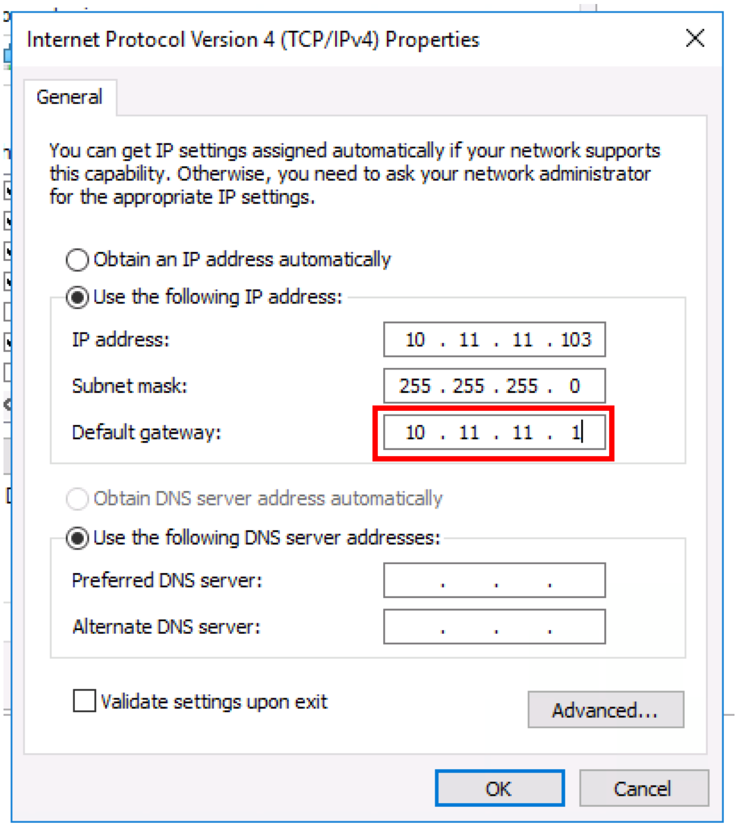

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 11 . 103

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 11 . 11 . 1

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 7.1-2: Internet Protocol Version 4 Properties PC-3

8. Ping the default gateway IP address. Ping should be successful.

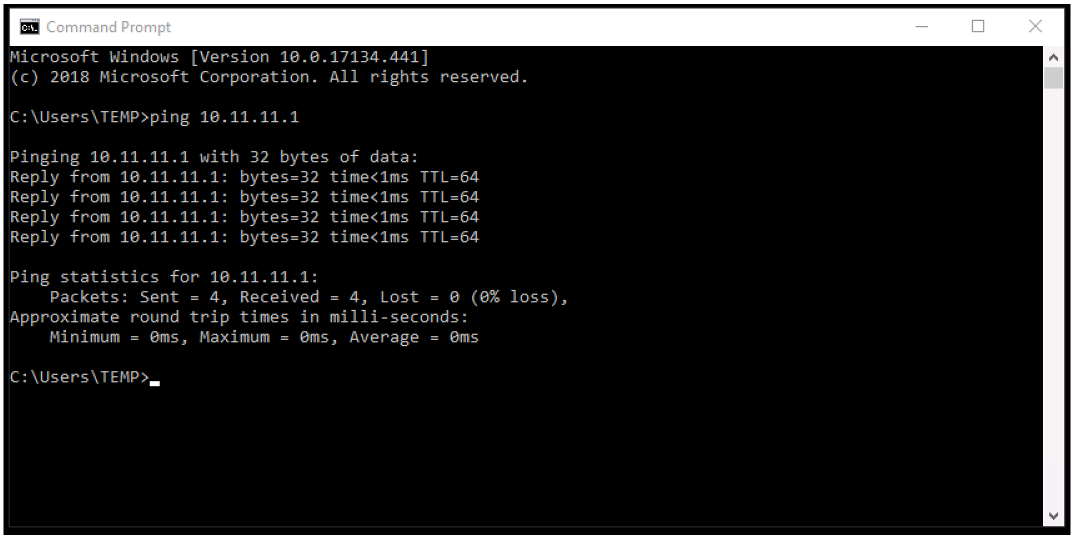

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\TEMP>ping 10.11.11.1

Pinging 10.11.11.1 with 32 bytes of data:
Reply from 10.11.11.1: bytes=32 time<1ms TTL=64
Reply from 10.11.11.1: bytes=32 time<1ms TTL=64
Reply from 10.11.11.1: bytes=32 time<1ms TTL=64
Reply from 10.11.11.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.11.11.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\TEMP>_
```

Figure 7.1-3: Ping to PC-3’s default gateway.

PC-4

9. Access PC-4
10. Repeat steps 7 and 8 using 10.X.12.1 instead.

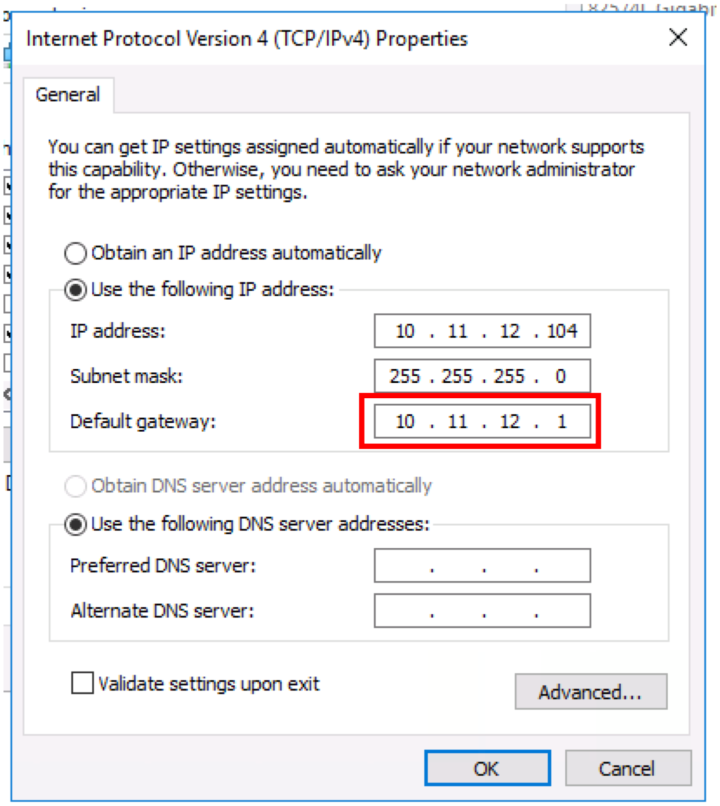

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 12 . 104

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 11 . 12 . 1

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 7.1-4: Internet Protocol Version 4 Properties PC-3

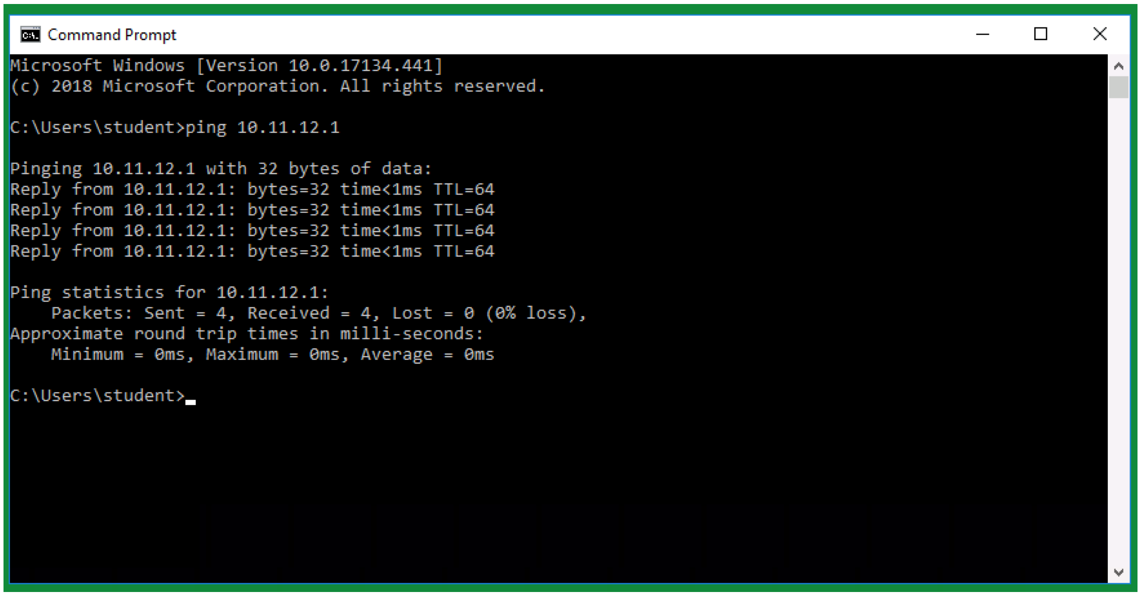

Command Prompt

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.12.1

Pinging 10.11.12.1 with 32 bytes of data:
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.11.12.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\student>_
```

Figure 7.1-5: Ping to PC-4’s default gateway.

11. From PC-4 ping PC-3 (10.X.11.103). Ping should be successful.

TIP: If ping is not successful, then it is possible that the Windows Firewall is dropping the ICMP traffic. Confirm the firewall is disabled on both PC-3 and PC-4.

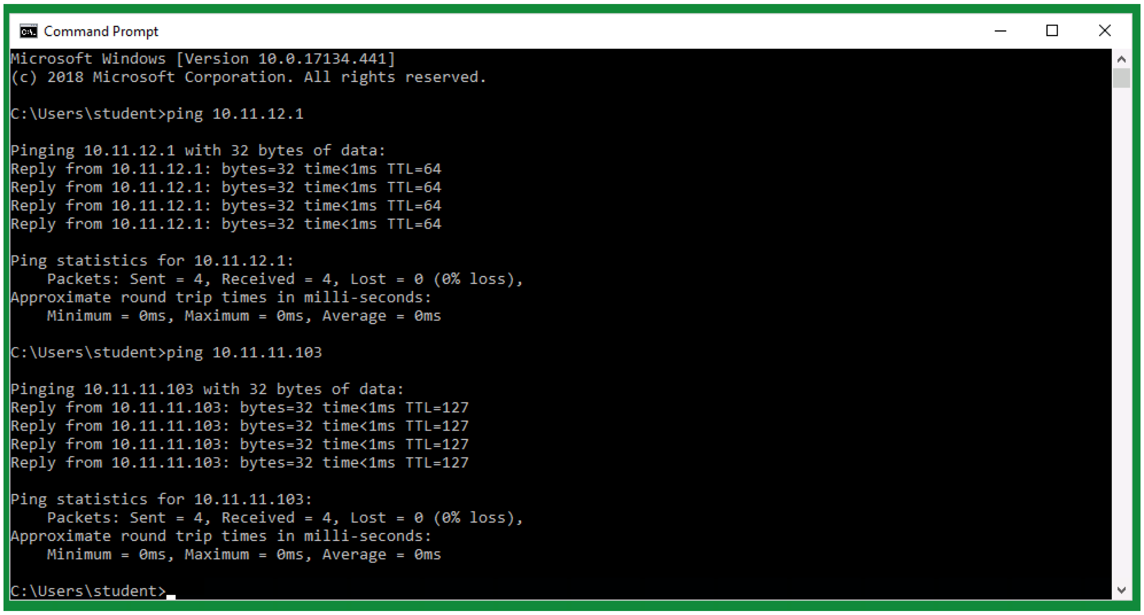

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.12.1

Pinging 10.11.12.1 with 32 bytes of data:
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64
Reply from 10.11.12.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.11.12.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\student>ping 10.11.11.103

Pinging 10.11.11.103 with 32 bytes of data:
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127

Ping statistics for 10.11.11.103:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\student>
```

Figure 7.1-6: Ping to PC-3.

Task 3: Explore End to End Packet Delivery

Objectives

In this part of the lab you will explore end to end packet delivery. You will examine Ethernet and IP headers, their addressing, and some of their fields using an open source traffic analysis tool called Wireshark. Wireshark will become an essential component of your networking troubleshooting tool kit.

Steps

Core-1 (via PC-1)

1. Open an SSH session to Core-1. Login using **cxfX/aruba123**
2. Clear ARP entries associated to PC-3 and PC-4 IP addresses (10.X.11.103 and 10.X.12.104 respectively).

```
Core-1#
Core-1# clear arp vrf default ipv4 10.X.11.103
Core-1# clear arp vrf default ipv4 10.X.12.104
```

PC-4

3. Access PC-4.
4. Right click the Command Prompt icon in the “**Start Bar**”, then right click the “**Command Prompt**” option that shows up or type in “**Cmd**” and select “**Run as Administrator**” in the menu that appears.

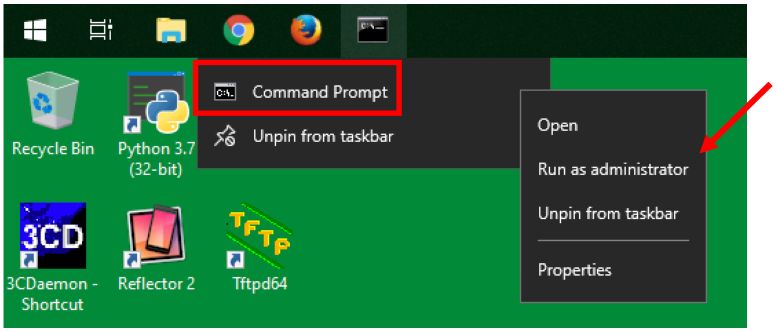

Figure 7.1-7: Command Prompt.

5. To accept the Windows warning below click on **yes**.

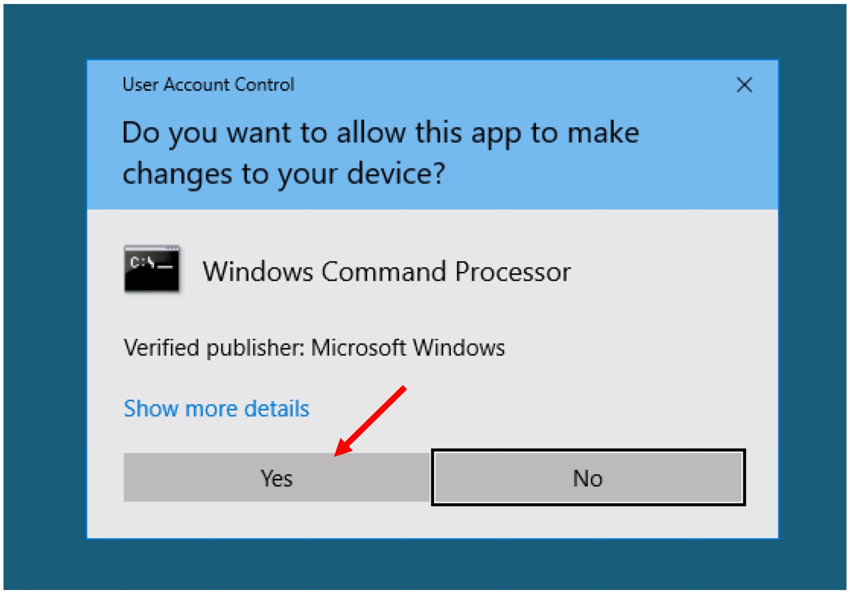
User Account Control

Do you want to allow this app to make changes to your device?

Windows Command Processor

Verified publisher: Microsoft Windows

Figure 7.1-8 Windows User Account Control.

6. Run the **arp -d** command to flush the ARP table of the host.

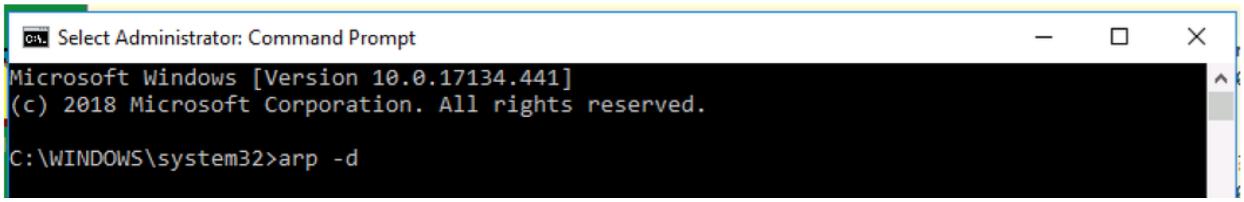

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>arp -d
```

Figure 7.1-9 Flush ARP table PC-4.

7. Run the **arp -a** command to display the ARP table in the host.

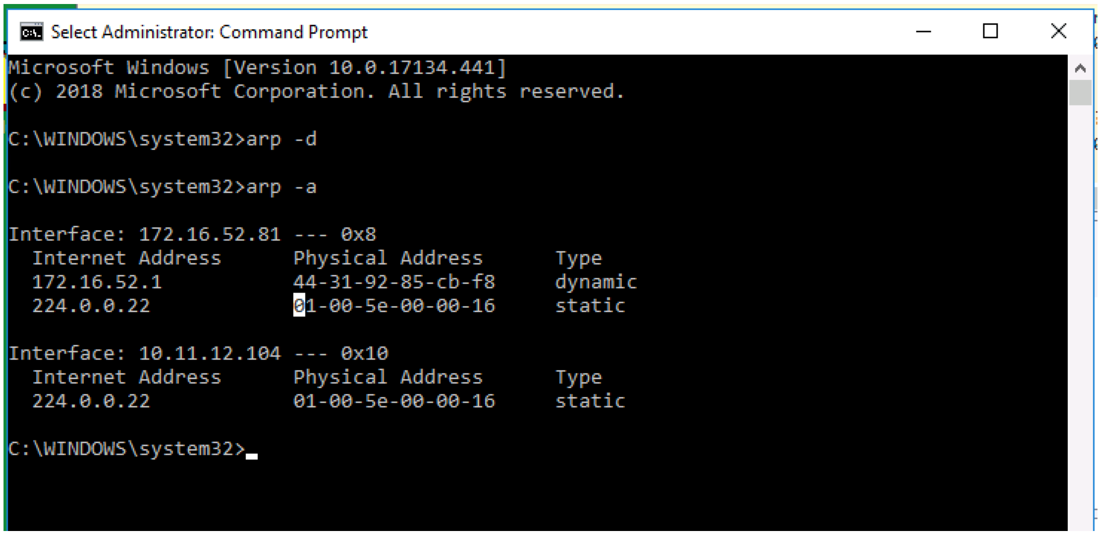

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>arp -d

C:\WINDOWS\system32>arp -a

Interface: 172.16.52.81 --- 0x8
  Internet Address      Physical Address      Type
  172.16.52.1           44-31-92-85-cb-f8     dynamic
  224.0.0.22            01-00-5e-00-00-16     static

Interface: 10.11.12.104 --- 0x10
  Internet Address      Physical Address      Type
  224.0.0.22            01-00-5e-00-00-16     static

C:\WINDOWS\system32>_
```

Figure 7.1-10 Display ARP table.

8. Open Wireshark from a shortcut on the Desktop.
9. Double click the “**Lab NIC**” entry. That will begin the packet capture on that interface. You will see gratuitous ARP messages coming from 10.X.12.1 (Core-1).

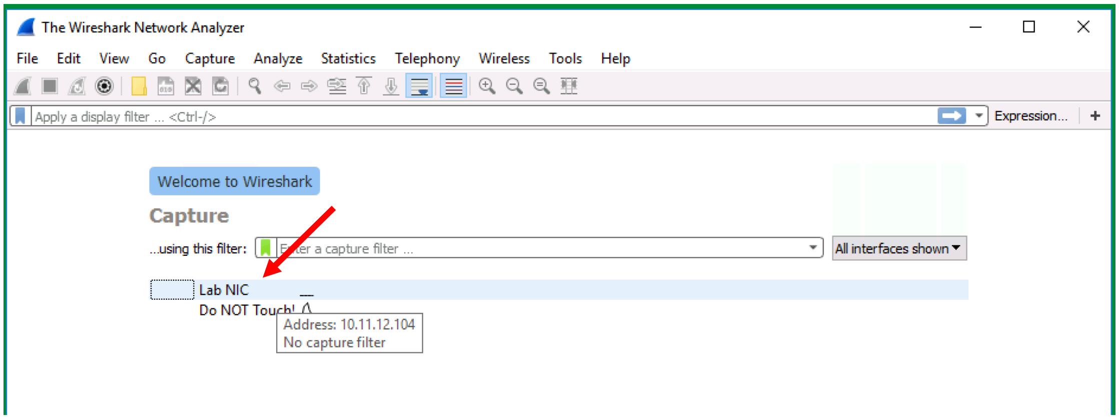

Lab NIC
Do NOT Touch!
Address: 10.11.12.104
No capture filter

Figure 7.1-11: Wireshark.

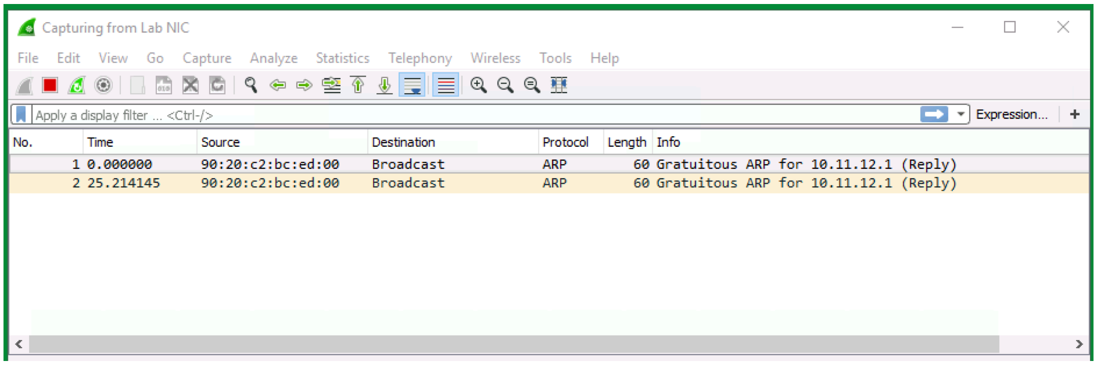
Capturing from Lab NIC
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Apply a display filter ... <Ctrl-/>
Expression...

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	90:20:c2:bc:ed:00	Broadcast	ARP	60	Gratuitous ARP for 10.11.12.1 (Reply)
2	25.214145	90:20:c2:bc:ed:00	Broadcast	ARP	60	Gratuitous ARP for 10.11.12.1 (Reply)

Figure 7.1-12: Gratuitous ARP.

IMPORTANT: Address Resolution Protocol (ARP) is a protocol that assists in IP Layer 3 to Ethernet/802.11 Layer 2 address resolution. When devices create an IP packet, they always have to find out the MAC address of the next hop (either the IP gateway when the Layer 3 destination is in a remote segment, or the destination host if it happens to be in the local segment of the sender). An IP packet cannot be sent out to the physical medium (copper, radio frequency or fiber) without a Data Link layer header. A Data Link layer header requires an address in order to be forwarded at layer 2 (e.g. Ethernet MAC, Frame Relay DCI, 802.11 BSSID, etc.).

IMPORTANT: AOS-CX advertises GARP packets every 25 seconds on the interfaces that have IP addresses. This updates any IP neighbor's ARP table and provides the resolution information in advance. However, operating systems like Microsoft Windows ignore these packets for security reasons.

10. In the filter, type "**(arp && not arp.isgratuitous) || ip.addr == 10.X.11.103**" with no quotes and hit **[Enter]**. That will instruct Wireshark to only display ARP non gratuitous messages and IP packets that include PC-3's IP address.

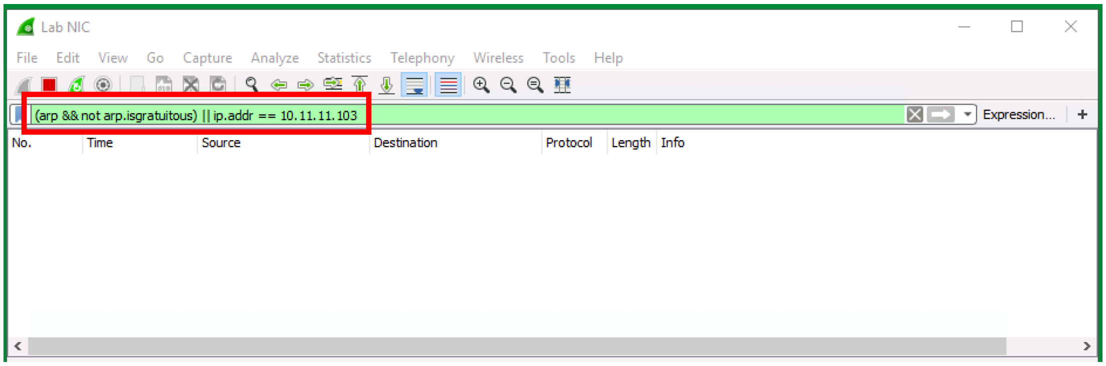
Lab NIC
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Expression… +
No. Time Source Destination Protocol Length Info

(arp && not arp.isgratuitous) || ip.addr == 10.11.11.103

Figure 7.1-13: Wireshark Filter.

PC-3

11. Move to PC-3.
12. Repeat steps 4 to 10 on PC-3 using 10.X.12.104 in the Wireshark filter.

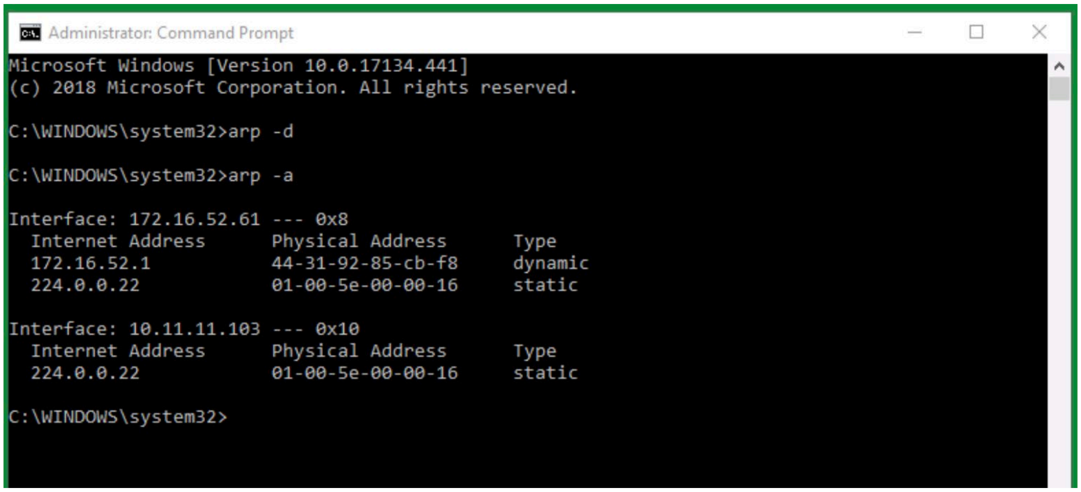

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>arp -d

C:\WINDOWS\system32>arp -a

Interface: 172.16.52.61 --- 0x8
  Internet Address      Physical Address      Type
  172.16.52.1           44-31-92-85-cb-f8     dynamic
  224.0.0.22            01-00-5e-00-00-16     static

Interface: 10.11.11.103 --- 0x10
  Internet Address      Physical Address      Type
  224.0.0.22            01-00-5e-00-00-16     static

C:\WINDOWS\system32>
```

Figure 7.1-14: Flush and display ARP table

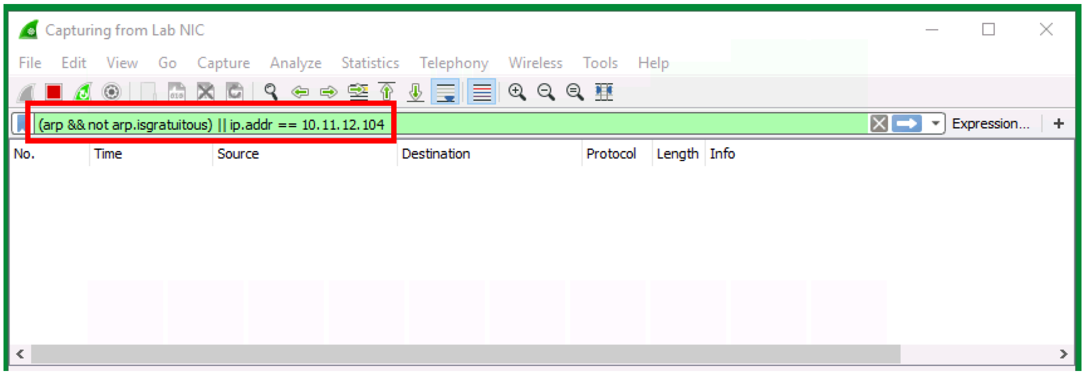
Capturing from Lab NIC
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Expression… +
No. Time Source Destination Protocol Length Info

(arp && not arp.isgratuitous) || ip.addr == 10.11.12.104

Figure 7.1-15: Wireshark Filter.

13. Run a custom ping on the command prompt using the following command: "**ping -n 1 10.X.12.104**" with no quotes. This command will trigger a single ICMP echo towards PC-4's IP address.

PC-3 and PC-4

14. Stop the Wireshark capture in both stations.

To begin the analysis, keep in mind what devices are involved in the packet forwarding. Use figure 7.1-16 as a reference.

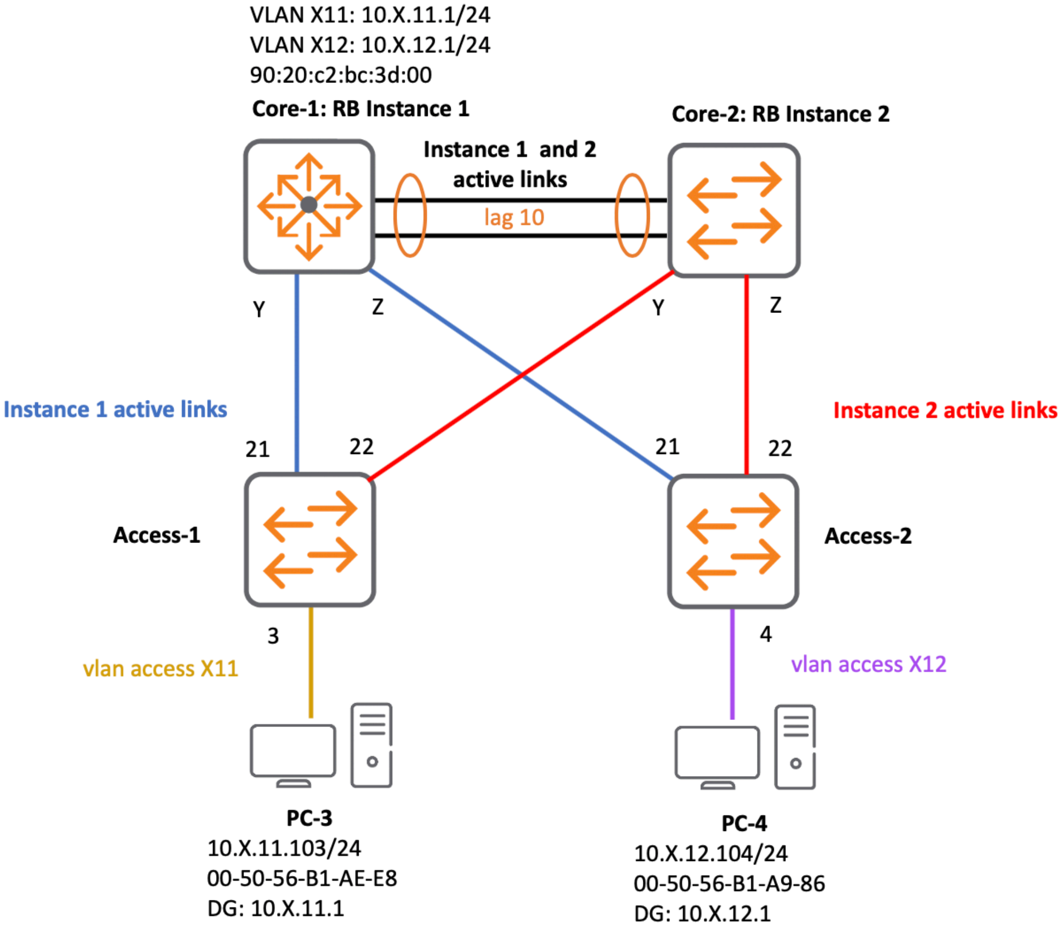

Figure 7.1-16: Topology for Traffic Analysis.

PC-3

In Wireshark you will see 6 frames in the capture, two of them are ICMP (pink packets), the four in yellow are ARP.

> **TIP:** Packets might be in a different order because there are limited resources assigned to client VMs. Nonetheless, the explanation below should help you know the order packets are sent.

15. Select the packet where its Destination equals “Broadcast”, that is an ARP request. Then look at the packet details section. You will see three gray rows, the first is the summary of the packet, the second is the Layer 2 header, and the third is the actual ARP payload.

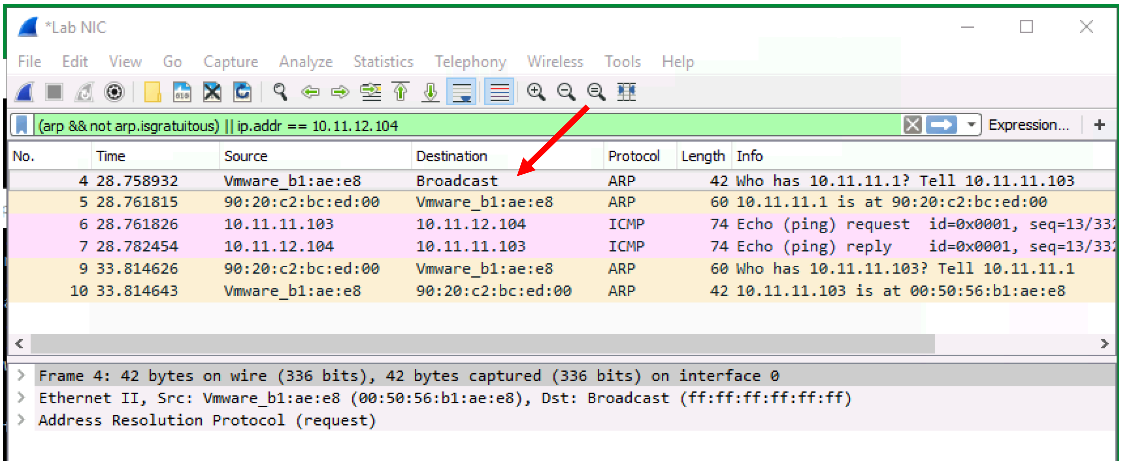

(arp && not arp.isgratuitous) || ip.addr == 10.11.12.104

```
No.  Time       Source             Destination        Protocol Length Info
   4 28.758932  Vmware_b1:ae:e8    Broadcast          ARP        42 Who has 10.11.11.1? Tell 10.11.11.103
   5 28.761815  90:20:c2:bc:ed:00  Vmware_b1:ae:e8    ARP        60 10.11.11.1 is at 90:20:c2:bc:ed:00
   6 28.761826  10.11.11.103       10.11.12.104       ICMP       74 Echo (ping) request  id=0x0001, seq=13/33
   7 28.782454  10.11.12.104       10.11.11.103       ICMP       74 Echo (ping) reply    id=0x0001, seq=13/33
   9 33.814626  90:20:c2:bc:ed:00  Vmware_b1:ae:e8    ARP        60 Who has 10.11.11.103? Tell 10.11.11.1
  10 33.814643  Vmware_b1:ae:e8    90:20:c2:bc:ed:00  ARP        42 10.11.11.103 is at 00:50:56:b1:ae:e8
```

```
> Frame 4: 42 bytes on wire (336 bits), 42 bytes captured (336 bits) on interface 0
> Ethernet II, Src: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
> Address Resolution Protocol (request)
```

Figure 7.1-17: Traffic Analysis 1

16. Select the Ethernet Layer 2 header and expand it.

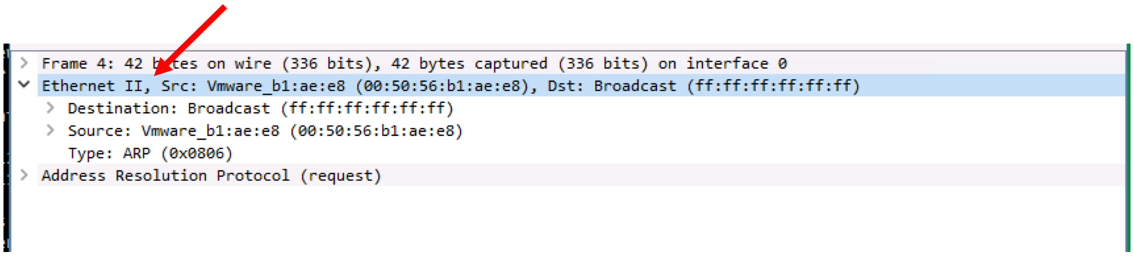

```
> Frame 4: 42 bytes on wire (336 bits), 42 bytes captured (336 bits) on interface 0
v Ethernet II, Src: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
  > Destination: Broadcast (ff:ff:ff:ff:ff:ff)
  > Source: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8)
    Type: ARP (0x0806)
> Address Resolution Protocol (request)
```

Figure 7.1-18: Traffic Analysis 2

What is the Destination MAC address?

____

What is the Source MAC address?

____

What is the Ethertype value?

____

ANSWER: The Destination MAC is all F's, which is the Broadcast MAC address, while the source is PC-3's MAC address. The Ethertype value is 0x0806 or ARP. This alerts the Layer 2 process what kind protocol or header comes next.

IMPORTANT: In Ethernet encapsulation, the destination MAC address is one of the first values in the packet. This helps the Layer 2 switch start the forwarding decision and processing of the frame as soon as it ingresses on the inbound port. This drastically enhances the throughput of the device.

17. Expand and select the third row (ARP Payload). This is an ARP request.

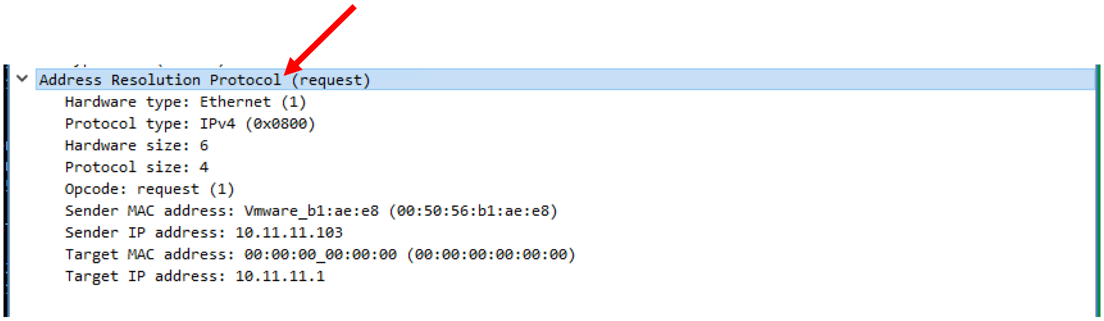

```
Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8)
    Sender IP address: 10.11.11.103
    Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
    Target IP address: 10.11.11.1
```

Figure 7.1-19: Traffic Analysis 3

What are the Sender MAC and IP addresses?

____

Who do they belong to?

____

What are the target MAC and IP addresses?

____

Why is the MAC address all 0’s?

____

____

What is the main purpose of this packet?

____

____

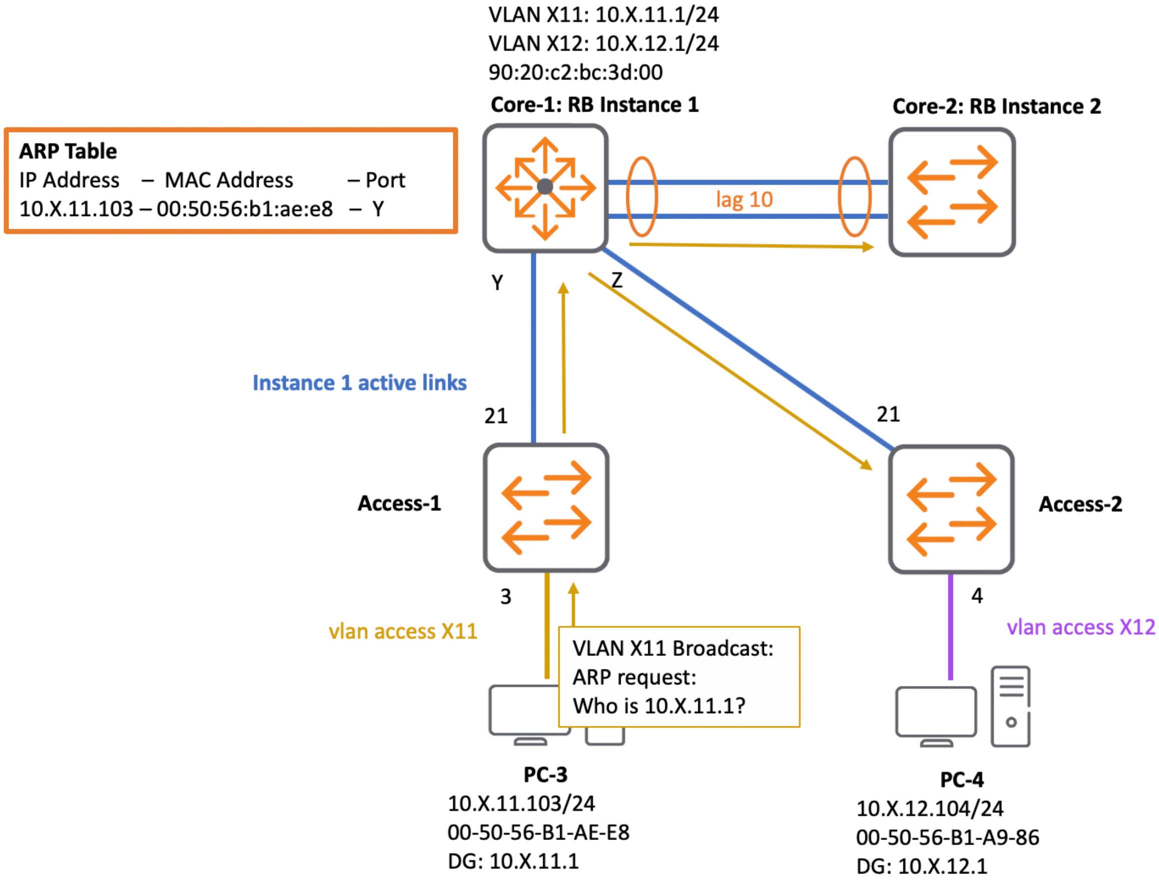

Figure 7.1-20: Topology for Traffic Analysis 2

NOTE: The destination of the packet is not a local segment (10.X.12.103), therefore PC-3 cannot reach it directly using Layer 2 but needs to send it to the default gateway (10.X.11.1). The default gateway will take the packet and route it out using Layer 3.

To do this, PC-3 has to take the ICMP echo request (from the ping command) and return it to Core-1 on VLAN X11. The IP header of the ICMP Echo request will remain untouched, however it has to be encapsulated with an Ethernet Layer 2 header to forward it.

In order to achieve this, PC-3 needs to know Core-1's MAC address so it can complete the Ethernet header generation. This process is known as Layer 3 to Layer 2 address resolution and requires ARP. Since you initially deleted PC-3's ARP table, it must send out an ARP request first, this packet uses the broadcast destination MAC address in order to assure it reaches all devices in the common VLAN.

When the broadcast is received by Access-1, it floods it across all ports in STP Forwarding mode for VLAN X11 except the sending port (port 3). Even though this is a broadcast packet, Access-2 does not decapsulate and process it beyond Layer 2 because the Ethertype 0x0806 tells the switch that an ARP packet will follow. Since ARP is an IP protocol (Layer 3) and Access-1 is not currently running Layer 3, there is no reason to keep inspecting the packet.

Core-1 receives the packet on port Y. Core-1 broadcasts the packet on all ports in Forwarding mode on VLAN X11 (Z and LAG 10). When the packet is received by Core-2 and Access-2 they just drop it.

When Core-1 looks at the Ethertype (ARP), it inspects the header at Layer 3 because IP is running on interface VLAN X11. After inspecting the ARP request, Core-1 recognizes the payload is asking for its own IP and prepares the reply.

18. Select the ARP reply (frame #5 in the figure below).

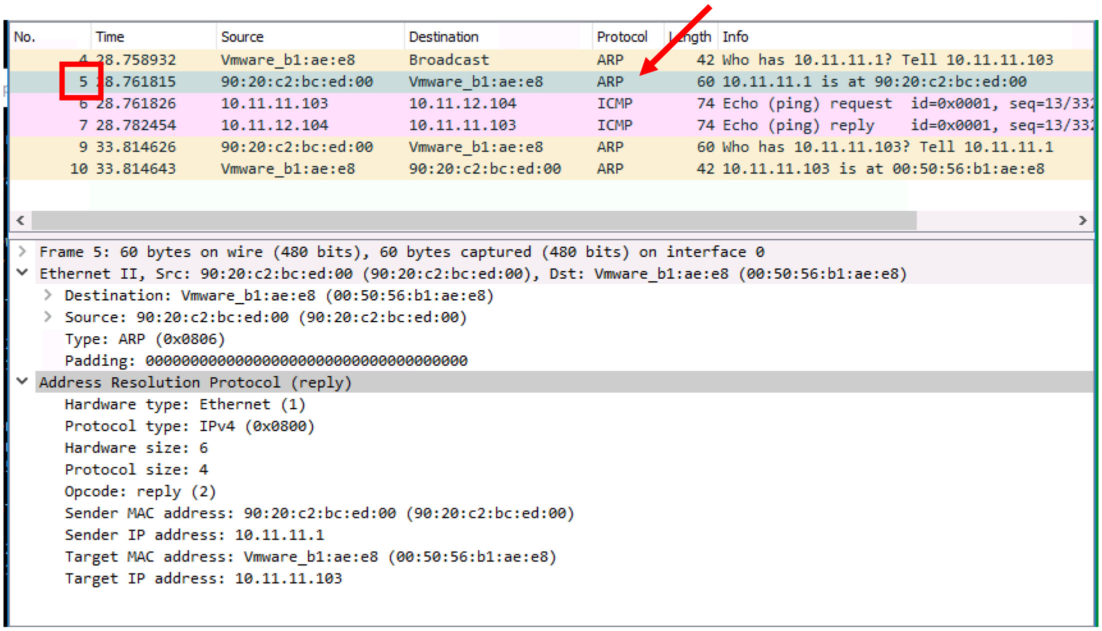

```
No.   Time        Source              Destination         Protocol  Length Info
    4 28.758932   Vmware_b1:ae:e8     Broadcast           ARP         42 Who has 10.11.11.1? Tell 10.11.11.103
    5 28.761815   90:20:c2:bc:ed:00   Vmware_b1:ae:e8     ARP         60 10.11.11.1 is at 90:20:c2:bc:ed:00
    6 28.761826   10.11.11.103        10.11.12.104        ICMP        74 Echo (ping) request  id=0x0001, seq=13/33
    7 28.782454   10.11.12.104        10.11.11.103        ICMP        74 Echo (ping) reply    id=0x0001, seq=13/33
    9 33.814626   90:20:c2:bc:ed:00   Vmware_b1:ae:e8     ARP         60 Who has 10.11.11.103? Tell 10.11.11.1
   10 33.814643   Vmware_b1:ae:e8     90:20:c2:bc:ed:00   ARP         42 10.11.11.103 is at 00:50:56:b1:ae:e8
```

```
> Frame 5: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
v Ethernet II, Src: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00), Dst: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8)
  > Destination: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8)
  > Source: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
    Type: ARP (0x0806)
    Padding: 000000000000000000000000000000000000
v Address Resolution Protocol (reply)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: reply (2)
    Sender MAC address: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
    Sender IP address: 10.11.11.1
    Target MAC address: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8)
    Target IP address: 10.11.11.103
```

Figure 7.1-21: Traffic Analysis 4

In the Ethernet header, what are the Destination and Source MAC addresses?

____

What kind of packet is this: Unicast, Broadcast or Multicast?

____

In the ARP header, what are the Sender MAC and IP addresses?

____

What are the Target MAC and IP addresses?

____

What is the main purpose of this packet?

____

____

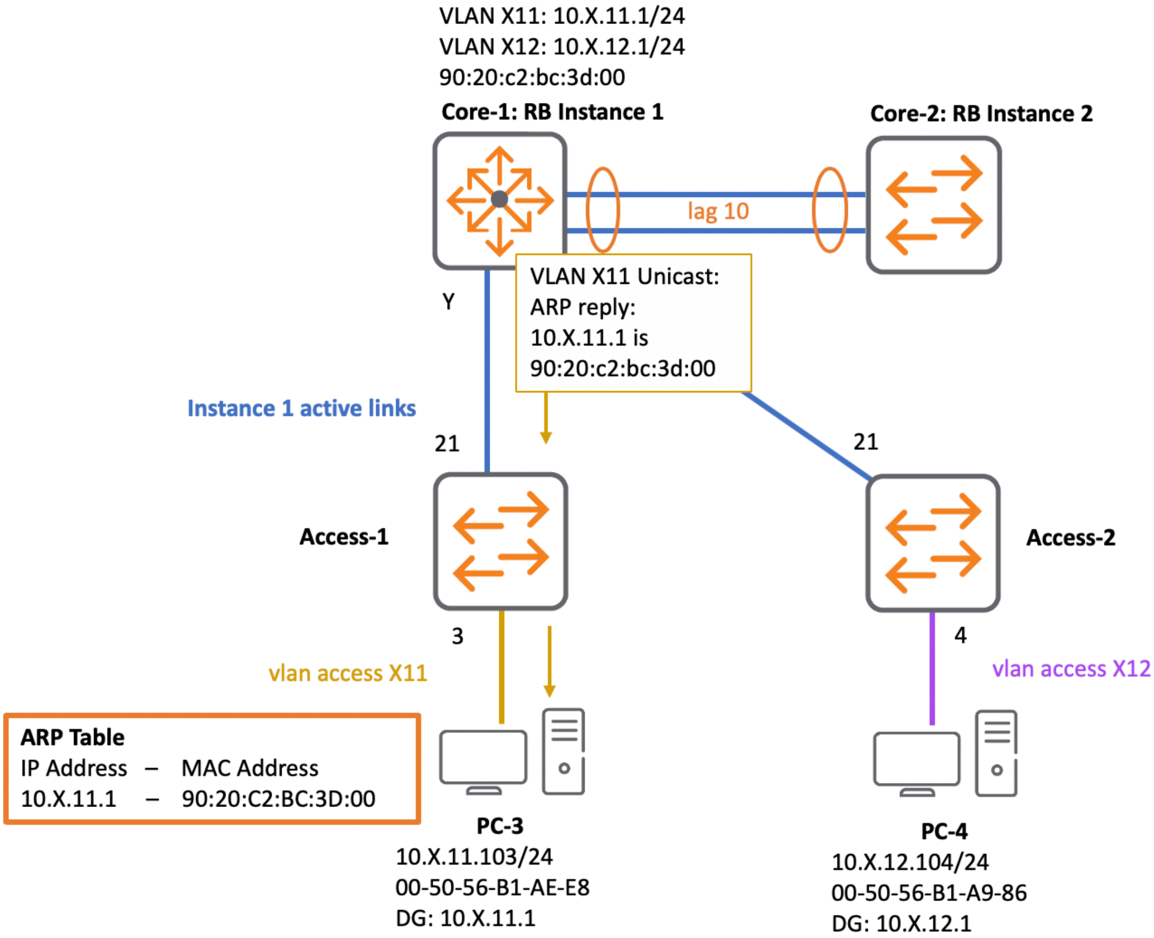

Figure 7.1-22: Topology for Traffic Analysis 3

NOTE: The Core-1 ARP reply is a regular unicast packet with the Layer 2 destination address of PC-3's MAC. The packet is received by Access-1. Access-1 uses its MAC Address table to forward the packet to port 3 and deliver it to PC-3.

When examining the Layer 3 payload, PC-3 recognizes this is the expected reply and uses the contents (Sender IP and MAC address) to generate an entry in its ARP table.

At this point PC-3 has completed the required Layer 2 to Layer 3 address resolution, now it can generate the Layer 2 header of the ICMP echo packet that it sends out.

19. Select the Echo (ping) request entry (frame #6 in the figure below), then expand the IP and ICMP headers.

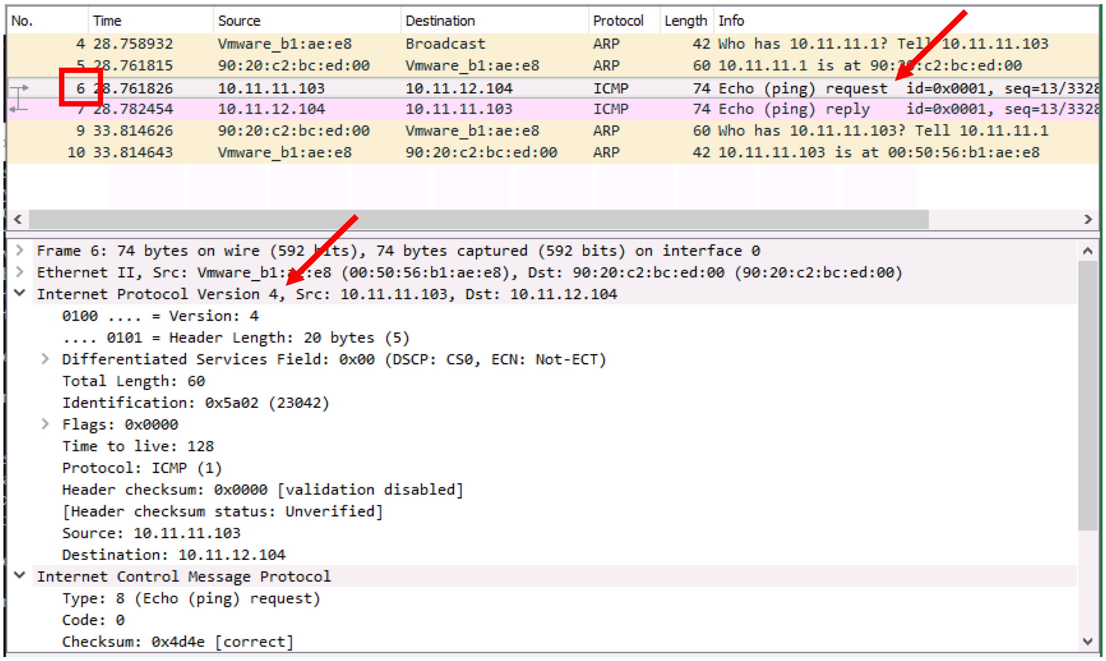

```
No.  Time       Source             Destination        Protocol  Length Info
   4 28.758932  Vmware_b1:ae:e8    Broadcast          ARP         42 Who has 10.11.11.1? Tell 10.11.11.103
   5 28.761815  90:20:c2:bc:ed:00  Vmware_b1:ae:e8    ARP         60 10.11.11.1 is at 90:20:c2:bc:ed:00
   6 28.761826  10.11.11.103       10.11.12.104       ICMP        74 Echo (ping) request  id=0x0001, seq=13/3328
   7 28.782454  10.11.12.104       10.11.11.103       ICMP        74 Echo (ping) reply    id=0x0001, seq=13/3328
   9 33.814626  90:20:c2:bc:ed:00  Vmware_b1:ae:e8    ARP         60 Who has 10.11.11.103? Tell 10.11.11.1
  10 33.814643  Vmware_b1:ae:e8    90:20:c2:bc:ed:00  ARP         42 10.11.11.103 is at 00:50:56:b1:ae:e8
```

```
> Frame 6: 74 bytes on wire (592 bits), 74 bytes captured (592 bits) on interface 0
> Ethernet II, Src: Vmware_b1:ae:e8 (00:50:56:b1:ae:e8), Dst: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
v Internet Protocol Version 4, Src: 10.11.11.103, Dst: 10.11.12.104
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 60
    Identification: 0x5a02 (23042)
  > Flags: 0x0000
    Time to live: 128
    Protocol: ICMP (1)
    Header checksum: 0x0000 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.11.11.103
    Destination: 10.11.12.104
v Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x4d4e [correct]
```

Figure 7.1-23: Traffic Analysis 4

On the Ethernet header, what is the Ethertype value?

____

What encapsulation is that?

____

What is the Layer 2 destination address?

____

What is the Layer 2 source address?

____

From the IP header, what is the Layer 3 source address?

____

What is the Layer 3 destination address?

____

Why are the Layer 2 and Layer 3 source addresses the same device, while the Layer 2 and Layer 3 destination addresses are different devices?

____

____

____

____

ANSWER: At the time the ICMP Echo request packet is generated, the Layer 3 destination address is the host you want to ping (PC-4). However, PC-4 it is not present in VLAN X11, so the packet has to be handed over to Core-1 (the default gateway of PC-3). This makes Core-1 the layer 2 destination of the frame.

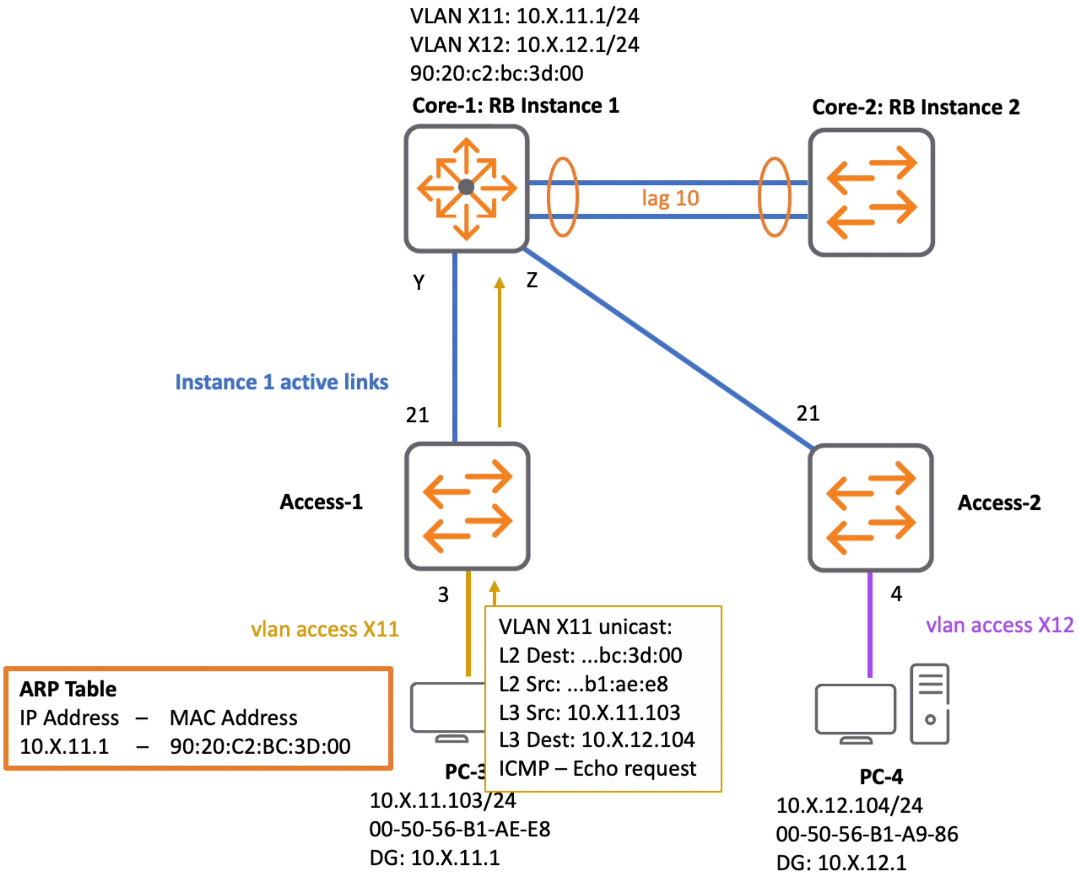

Figure 7.1-24: Topology for Traffic Analysis 4

What is the Time to Live value?

____

ANSWER: Time to Live is the maximum number of Layer 3 boundaries the packet will be able to cross before getting dropped.

What is the Protocol value?

____

NOTE: As mentioned in Module 1, IP protocol is used to signal the next layer protocol.

The following part of the process takes place on VLAN X12. Since PC-3 is not part of that broadcast domain, move to PC-4 and continue the packet analysis from there.

PC-4

20. Move to **PC-4**.
21. In Wireshark, Select the packet where its Destination equals "**Broadcast**" and expand the Address Resolution Protocol row in the packet details section.

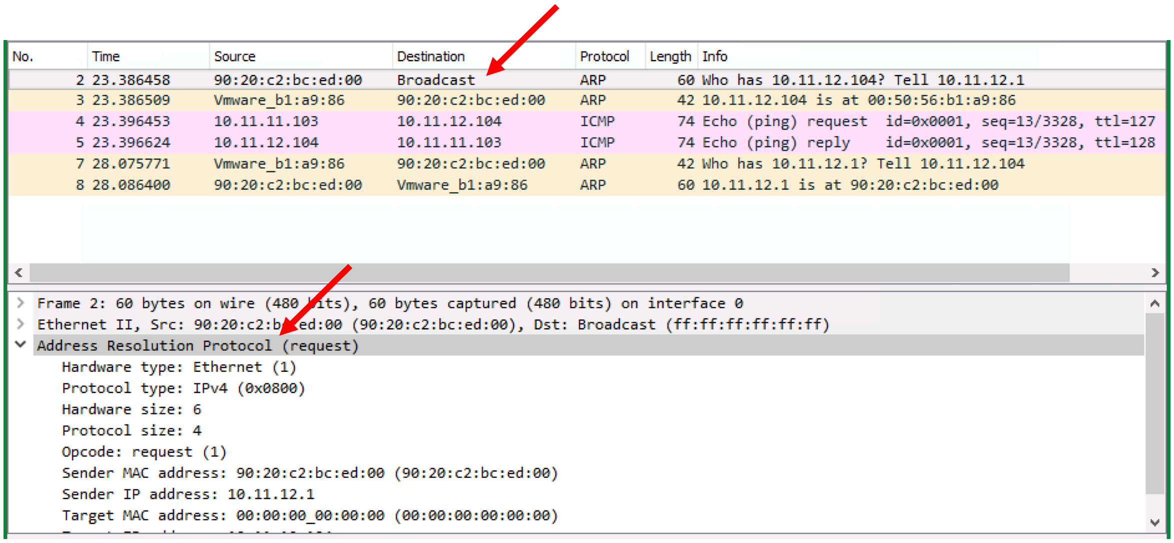

```
No.   Time       Source             Destination        Protocol Length Info
    2 23.386458  90:20:c2:bc:ed:00  Broadcast          ARP        60 Who has 10.11.12.104? Tell 10.11.12.1
    3 23.386509  Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP        42 10.11.12.104 is at 00:50:56:b1:a9:86
    4 23.396453  10.11.11.103       10.11.12.104       ICMP       74 Echo (ping) request  id=0x0001, seq=13/3328, ttl=127
    5 23.396624  10.11.12.104       10.11.11.103       ICMP       74 Echo (ping) reply    id=0x0001, seq=13/3328, ttl=128
    7 28.075771  Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP        42 Who has 10.11.12.1? Tell 10.11.12.104
    8 28.086400  90:20:c2:bc:ed:00  Vmware_b1:a9:86    ARP        60 10.11.12.1 is at 90:20:c2:bc:ed:00
```

```
> Frame 2: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
> Ethernet II, Src: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
v Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
    Sender IP address: 10.11.12.1
    Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
```

Figure 7.1-25: Traffic Analysis 5

On the ARP header, what are the Sender MAC and IP addresses?

____

What do they belong to?

____

What are the target MAC and IP addresses?

____

What is the main purpose of this packet?

____

____

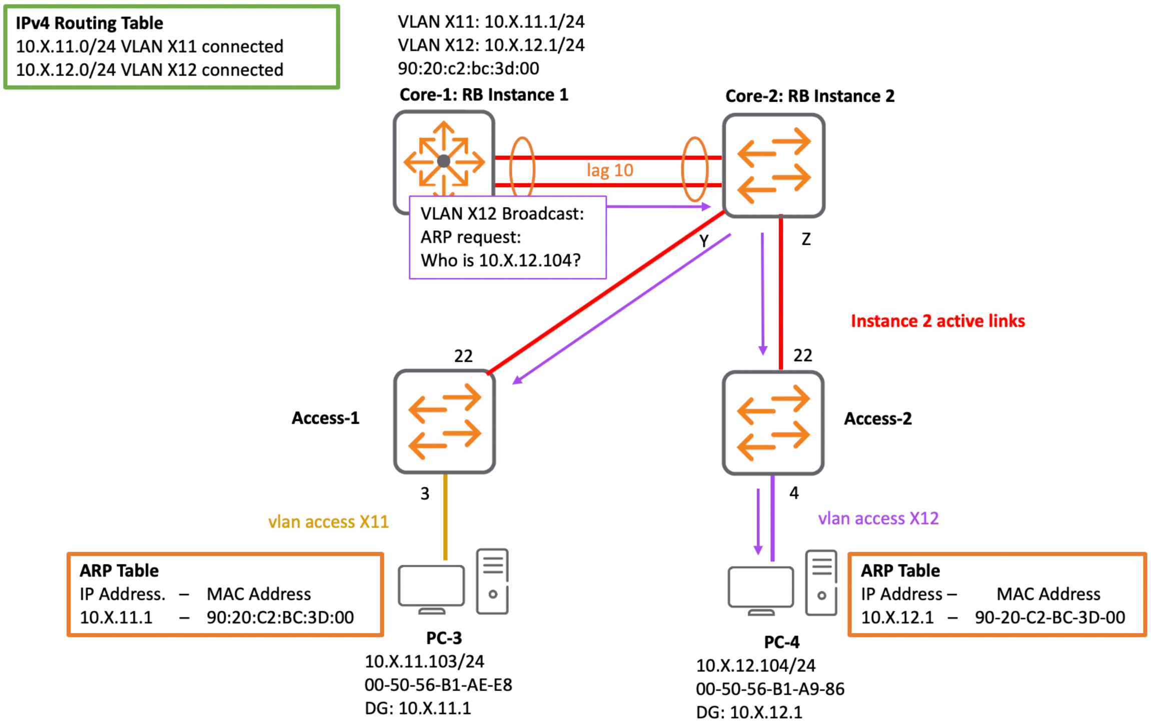

Figure 7.1-26: Topology for Traffic Analysis 5

NOTE: When Core-1 received the ICMP packet and decapsulated it up to Layer 3, it looked into the destination IPv4 address. Core-1 determines it is not the IP destination of the packet and must move the packet between VLANs (Inter-VLAN routing).

To route between VLANs, Core-1 examines its routing table. It looks for an entry with an IP prefix or network that includes the destination IP address. If several entries are found, then the longest match (the more specific route) is used. In the current routing table, there is a valid entry: 10.X.12.0/24 out of VLAN X12 that Core-1 can use. It is a connected route.

Core-1 is now like PC-3 at the beginning of the process. It knows which outbound Layer 3 interface to use but it must create the Layer 2 header, therefore it needs to perform another Layer 2/Layer 3 address resolution requesting PC-4's MAC address.

Core-1 creates the ARP request where the Target IP address is 10.X.12.104 and sends it as a broadcast flood out LAG 10 tagged in VLAN X12. Core-2 gets a copy and sends it out of port Y and Z, the packet dies on Access-1, but is retained in Access-2 who delivers it to PC-4.

22. Select the **ARP reply** from PC-4 to Core-2 (frame #3 in figure below).

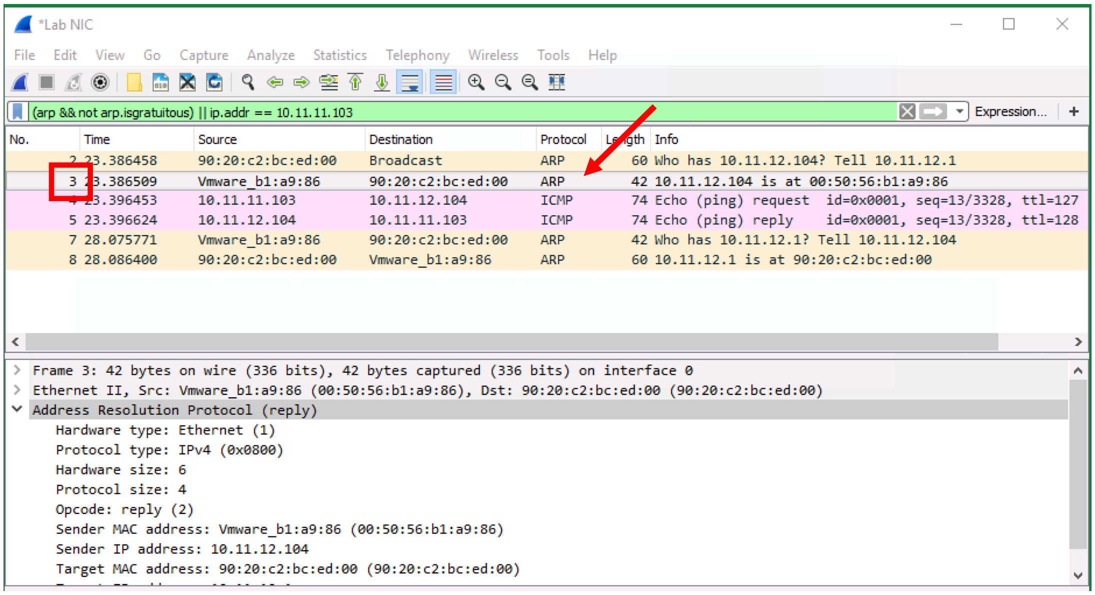
*Lab NIC
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Expression...
+

(arp && not arp.isgratuitous) || ip.addr == 10.11.11.103

```
No.     Time       Source             Destination        Protocol  Length Info
      2 23.386458  90:20:c2:bc:ed:00  Broadcast          ARP           60 Who has 10.11.12.104? Tell 10.11.12.1
      3 23.386509  Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP           42 10.11.12.104 is at 00:50:56:b1:a9:86
      4 23.396453  10.11.11.103       10.11.12.104       ICMP          74 Echo (ping) request  id=0x0001, seq=13/3328, ttl=127
      5 23.396624  10.11.12.104       10.11.11.103       ICMP          74 Echo (ping) reply    id=0x0001, seq=13/3328, ttl=128
      7 28.075771  Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP           42 Who has 10.11.12.1? Tell 10.11.12.104
      8 28.086400  90:20:c2:bc:ed:00  Vmware_b1:a9:86    ARP           60 10.11.12.1 is at 90:20:c2:bc:ed:00
```

```
> Frame 3: 42 bytes on wire (336 bits), 42 bytes captured (336 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a9:86 (00:50:56:b1:a9:86), Dst: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
v Address Resolution Protocol (reply)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: reply (2)
    Sender MAC address: Vmware_b1:a9:86 (00:50:56:b1:a9:86)
    Sender IP address: 10.11.12.104
    Target MAC address: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
```

Figure 7.1-27: Traffic Analysis 6

What is the source MAC address?

____

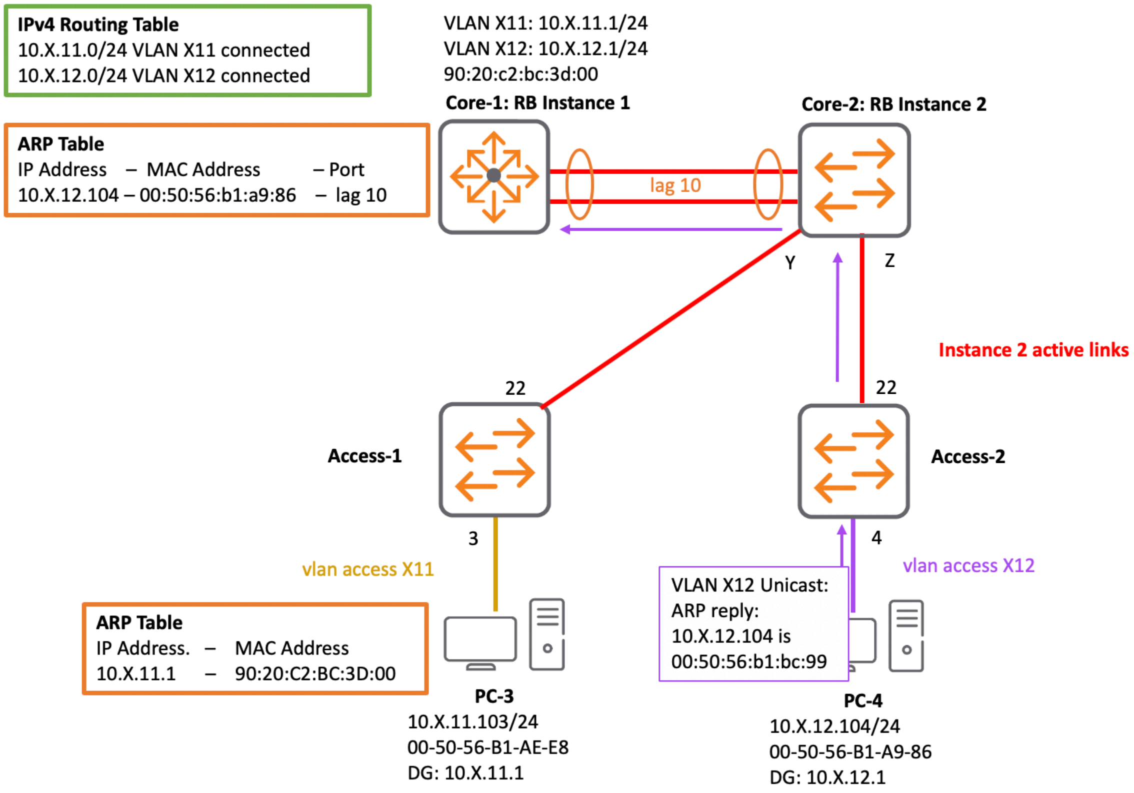

IPv4 Routing Table
10.X.11.0/24 VLAN X11 connected
10.X.12.0/24 VLAN X12 connected

ARP Table
IP Address – MAC Address – Port
10.X.12.104 – 00:50:56:b1:a9:86 – lag 10

ARP Table
IP Address. – MAC Address
10.X.11.1 – 90:20:C2:BC:3D:00

VLAN X12 Unicast:
ARP reply:
10.X.12.104 is
00:50:56:b1:bc:99

Figure 7.1-28: Topology for Traffic Analysis 6

NOTE: When PC-4 generates the ARP reply, this goes to Core-1. Core-1 updates its ARP table and is ready to deliver the ICMP echo message.

23. Select the ICMP echo message (frame #4 in the figure below). And focus on the Layer 2 and Layer 3 addresses.

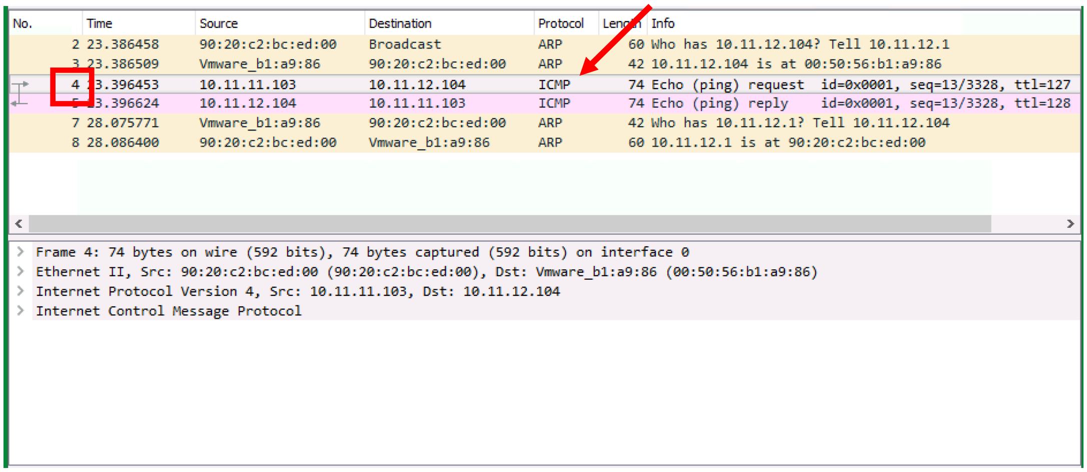

```
No.   Time        Source              Destination         Protocol Length Info
    2 23.386458   90:20:c2:bc:ed:00   Broadcast           ARP        60 Who has 10.11.12.104? Tell 10.11.12.1
    3 23.386509   Vmware_b1:a9:86     90:20:c2:bc:ed:00   ARP        42 10.11.12.104 is at 00:50:56:b1:a9:86
    4 23.396453   10.11.11.103        10.11.12.104        ICMP       74 Echo (ping) request  id=0x0001, seq=13/3328, ttl=127
    5 23.396624   10.11.12.104        10.11.11.103        ICMP       74 Echo (ping) reply    id=0x0001, seq=13/3328, ttl=128
    7 28.075771   Vmware_b1:a9:86     90:20:c2:bc:ed:00   ARP        42 Who has 10.11.12.1? Tell 10.11.12.104
    8 28.086400   90:20:c2:bc:ed:00   Vmware_b1:a9:86     ARP        60 10.11.12.1 is at 90:20:c2:bc:ed:00

> Frame 4: 74 bytes on wire (592 bits), 74 bytes captured (592 bits) on interface 0
> Ethernet II, Src: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00), Dst: Vmware_b1:a9:86 (00:50:56:b1:a9:86)
> Internet Protocol Version 4, Src: 10.11.11.103, Dst: 10.11.12.104
> Internet Control Message Protocol
```

Figure 7.1-29: Traffic Analysis 7

What are the Layer 2 destination and source addresses?

How did they change from step 18?

What are the Layer 3 destination and source addresses?

Did they change from step 18?

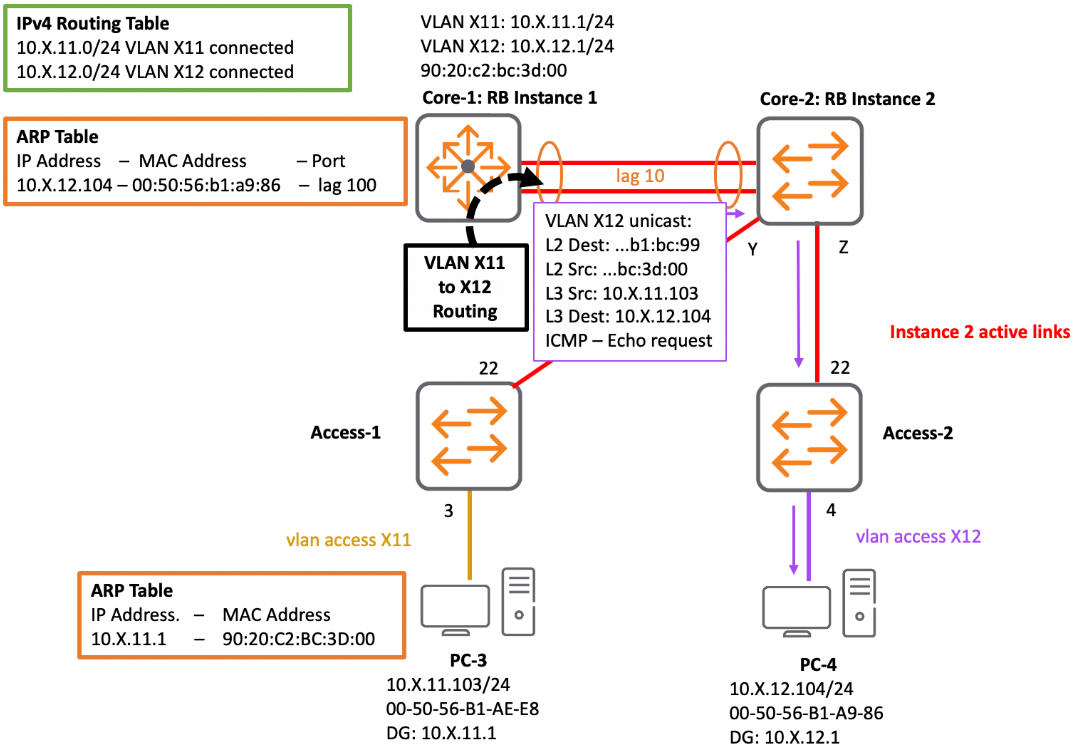

IPv4 Routing Table
10.X.11.0/24 VLAN X11 connected
10.X.12.0/24 VLAN X12 connected

ARP Table

IP Address	MAC Address	Port
10.X.12.104	00:50:56:b1:a9:86	lag 100

ARP Table

IP Address.	MAC Address
10.X.11.1	90:20:C2:BC:3D:00

VLAN X12 unicast:
L2 Dest: ...b1:bc:99
L2 Src: ...bc:3d:00
L3 Src: 10.X.11.103
L3 Dest: 10.X.12.104
ICMP – Echo request

Figure 7.1-30: Topology for Traffic Analysis 7

NOTE: After creating the Layer 2 header with PC-4's MAC address and looking into its MAC address table, Core-1 is ready to forward the packet using LAG 10 as the outbound interface for the unicast packet. When leaving Core-1 the packet crosses Core-2, Access-2 and finally gets to PC-4.

This new version of the packet has the Core-1 MAC address as its Layer 2 source address rather than its destination address (as it was in step 18) and PC-4 is now the new destination address. Layer 2 addresses change at each routing hop.

24. Select the second ARP request (frame #7 in the figure below) and inspect its contents.

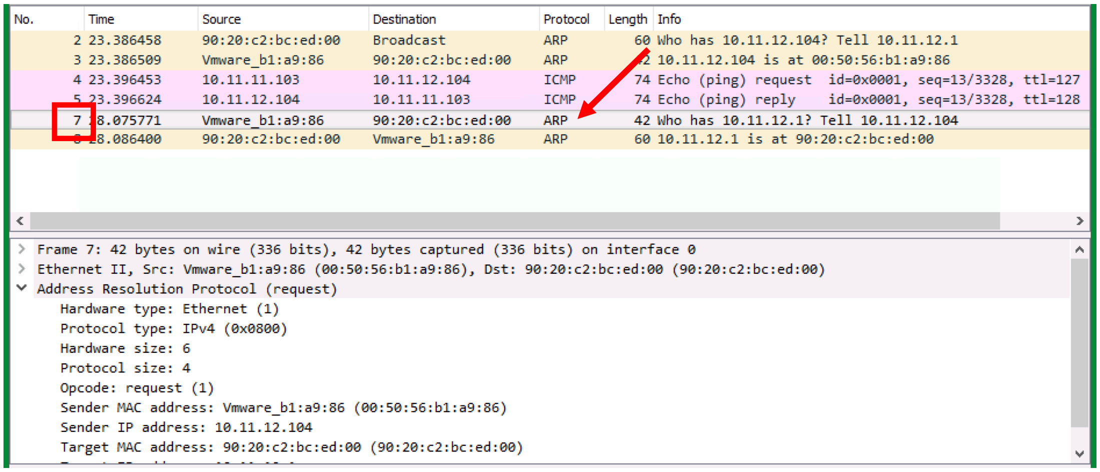

```
No.     Time        Source              Destination         Protocol  Length  Info
      2 23.386458   90:20:c2:bc:ed:00   Broadcast           ARP       60 Who has 10.11.12.104? Tell 10.11.12.1
      3 23.386509   Vmware_b1:a9:86     90:20:c2:bc:ed:00   ARP       42 10.11.12.104 is at 00:50:56:b1:a9:86
      4 23.396453   10.11.11.103        10.11.12.104        ICMP      74 Echo (ping) request  id=0x0001, seq=13/3328, ttl=127
      5 23.396624   10.11.12.104        10.11.11.103        ICMP      74 Echo (ping) reply    id=0x0001, seq=13/3328, ttl=128
      7 28.075771   Vmware_b1:a9:86     90:20:c2:bc:ed:00   ARP       42 Who has 10.11.12.1? Tell 10.11.12.104
      8 28.086400   90:20:c2:bc:ed:00   Vmware_b1:a9:86     ARP       60 10.11.12.1 is at 90:20:c2:bc:ed:00

> Frame 7: 42 bytes on wire (336 bits), 42 bytes captured (336 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a9:86 (00:50:56:b1:a9:86), Dst: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
v Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: Vmware_b1:a9:86 (00:50:56:b1:a9:86)
    Sender IP address: 10.11.12.104
    Target MAC address: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
```

Figure 7.1-31: Traffic Analysis 8

NOTE: Before replying, PC-4 (as Core-1 and PC-3 before it) needs to add its gateway MAC address to its ARP table. That triggers the ARP request seen in image above. In entry number 8, PC-4 gets an ARP reply back from Core-1.

25. Select the ICMP (ping) reply (frame #5 in the figure below).

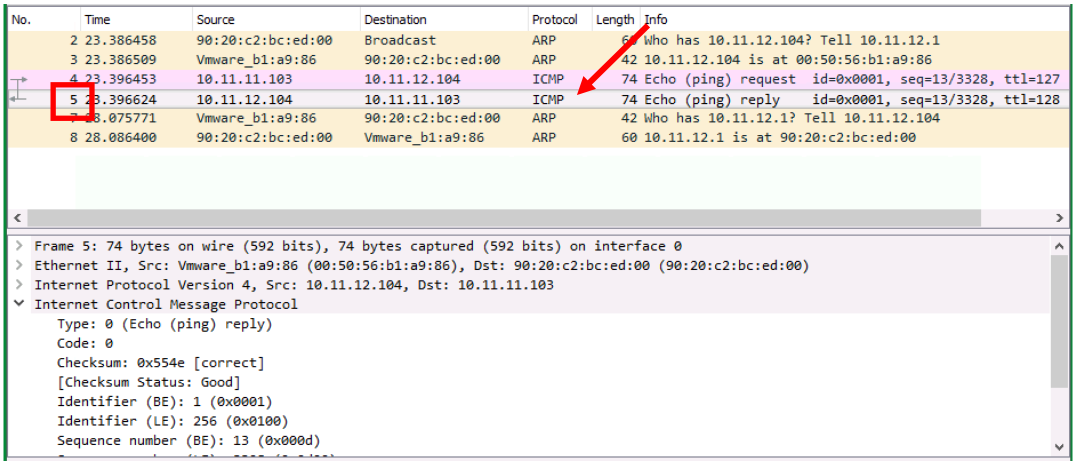

```
No.    Time         Source             Destination        Protocol  Length Info
     2 23.386458    90:20:c2:bc:ed:00  Broadcast          ARP       60 Who has 10.11.12.104? Tell 10.11.12.1
     3 23.386509    Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP       42 10.11.12.104 is at 00:50:56:b1:a9:86
     4 23.396453    10.11.11.103       10.11.12.104       ICMP      74 Echo (ping) request  id=0x0001, seq=13/3328, ttl=127
     5 23.396624    10.11.12.104       10.11.11.103       ICMP      74 Echo (ping) reply    id=0x0001, seq=13/3328, ttl=128
     7 28.075771    Vmware_b1:a9:86    90:20:c2:bc:ed:00  ARP       42 Who has 10.11.12.1? Tell 10.11.12.104
     8 28.086400    90:20:c2:bc:ed:00  Vmware_b1:a9:86    ARP       60 10.11.12.1 is at 90:20:c2:bc:ed:00

> Frame 5: 74 bytes on wire (592 bits), 74 bytes captured (592 bits) on interface 0
> Ethernet II, Src: Vmware_b1:a9:86 (00:50:56:b1:a9:86), Dst: 90:20:c2:bc:ed:00 (90:20:c2:bc:ed:00)
> Internet Protocol Version 4, Src: 10.11.12.104, Dst: 10.11.11.103
v Internet Control Message Protocol
    Type: 0 (Echo (ping) reply)
    Code: 0
    Checksum: 0x554e [correct]
    [Checksum Status: Good]
    Identifier (BE): 1 (0x0001)
    Identifier (LE): 256 (0x0100)
    Sequence number (BE): 13 (0x000d)
```

Figure 7.1-32: Traffic Analysis 9

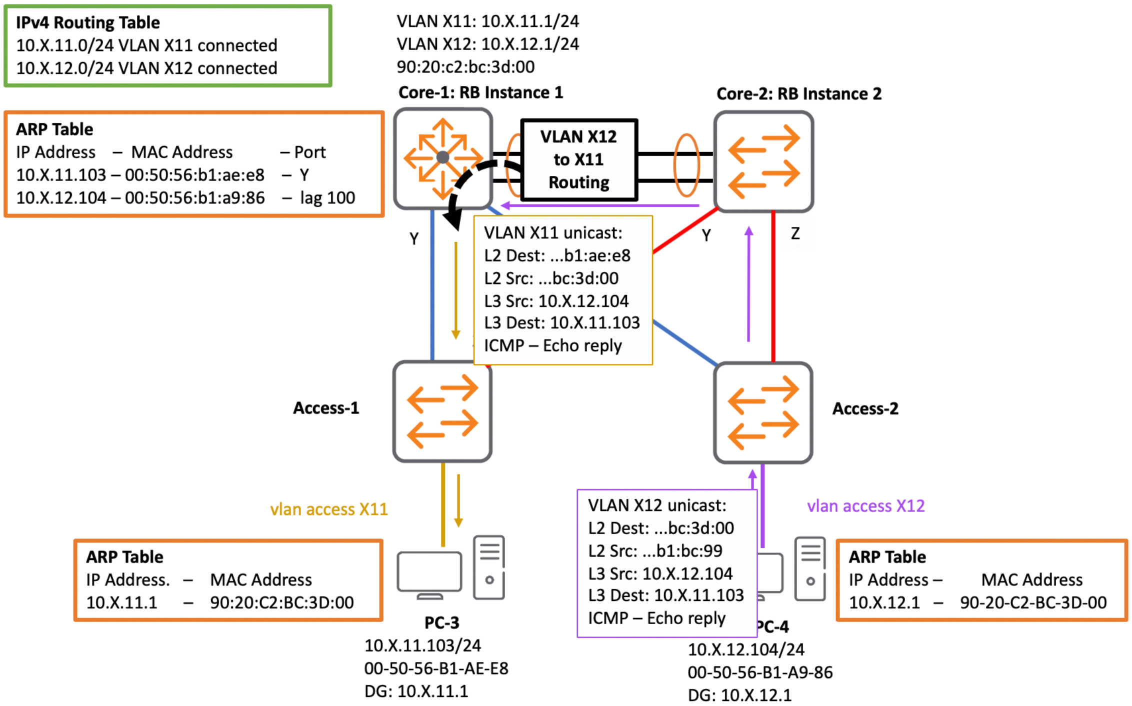

Figure 7.1-33: Topology for Traffic Analysis 8

NOTE: When PC-4, completes the encapsulation step, it sends the packet to Core-1. Again Core-1 has to perform an ARP lookup to add the PC-3 MAC address. After encapsulating the packet, Core-1 forwards the ICMP echo reply to PC-3 and the process ends.

26. Close Wireshark in both PCs.

Task 3: Save Your Configurations

Objectives

You will now proceed to save your configuration

Steps

Core-1 (via PC-1)

1. Save the current Core-1’s configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

You have completed Lab 7.1!

AOS-CX Switching Fundamentals

Lab 7.2: Creating a VRF

Overview

A few days after enabling routing in Core-1, BigStartup was notified that other tenants will also be connecting to the 8325-switch pair. Therefore, during a maintenance window, you will have to create a custom VRF for keeping local segments private and avoid traffic leaking.

Objectives

After completing this lab, you will be able to:

- Create a custom VRF
- Assign SVIs to VRF
- Explore VRF specific routing table

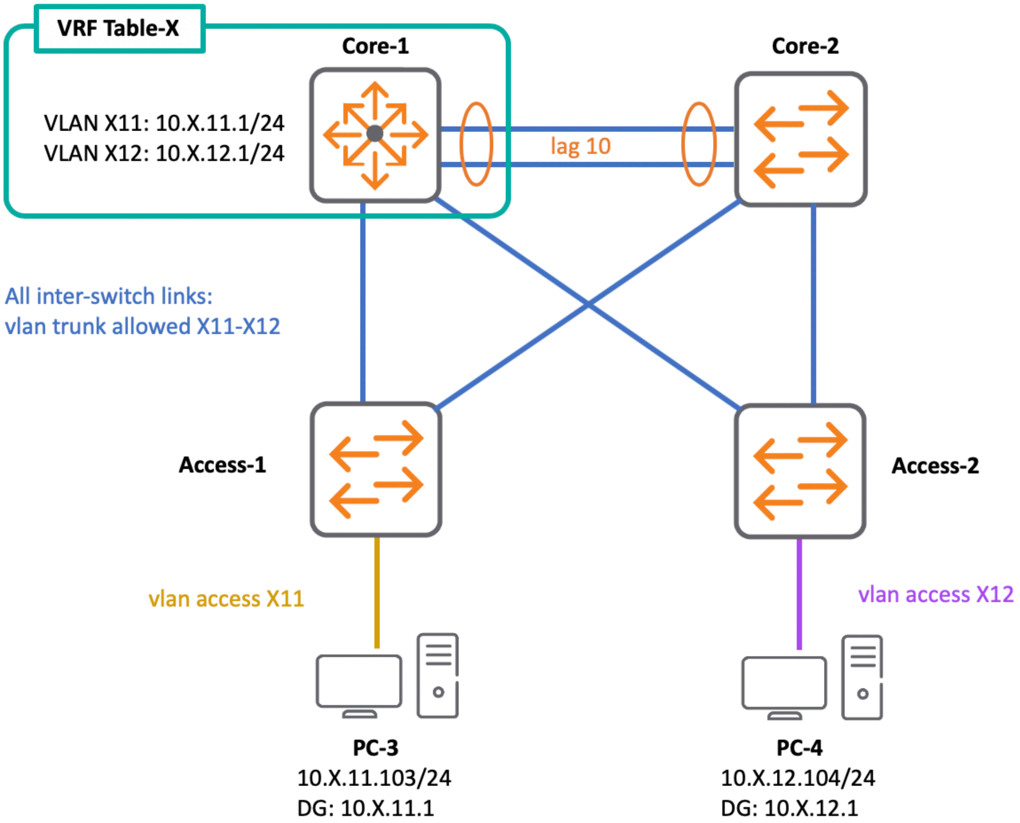

Figure 7.2-1: Lab Topology

Task 1: Create Table VRF

Objectives

In this step you will migrate your customer's network into an exclusive VRF. This requires creating it, assigning Layer 3 interfaces, and re-configuring the IP settings. Since the process might suspend IP services, a one-hour maintenance window has been scheduled for this task. You must act promptly!

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1.
2. Ping PC-3 (**10.X.11.103**) and PC-4 (**10.X.12.104**). Pings should be successful.

NOTE: Replace the highlighted "X" for your student table number.

```
Core-1# ping 10.X.11.103
PING 10.11.11.103 (10.11.11.103) 100(128) bytes of data.
108 bytes from 10.11.11.103: icmp_seq=1 ttl=128 time=19.0 ms
108 bytes from 10.11.11.103: icmp_seq=2 ttl=128 time=0.802 ms
108 bytes from 10.11.11.103: icmp_seq=3 ttl=128 time=0.642 ms
108 bytes from 10.11.11.103: icmp_seq=4 ttl=128 time=0.715 ms
108 bytes from 10.11.11.103: icmp_seq=5 ttl=128 time=0.706 ms

--- 10.11.11.103 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4051ms
rtt min/avg/max/mdev = 0.642/4.374/19.007/7.316 ms
Core-1#
```

```
Core-1# ping 10.X.12.104
PING 10.11.12.104 (10.11.12.104) 100(128) bytes of data.
108 bytes from 10.11.12.104: icmp_seq=1 ttl=128 time=20.2 ms
108 bytes from 10.11.12.104: icmp_seq=2 ttl=128 time=0.817 ms
108 bytes from 10.11.12.104: icmp_seq=3 ttl=128 time=0.831 ms
108 bytes from 10.11.12.104: icmp_seq=4 ttl=128 time=0.762 ms
108 bytes from 10.11.12.104: icmp_seq=5 ttl=128 time=0.682 ms

--- 10.11.12.104 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4053ms
rtt min/avg/max/mdev = 0.682/4.661/20.217/7.778 ms
```

```
Core-1#
```

3. Move to configuration mode and create a VRF named **TABLE-X**.

NOTE: Replace the highlighted “X” for your student table number.

```
Core-1# configure terminal
Core-1(config)# vrf TABLE-X
Core-1(config-vrf)# exit
Core-1(config)#
```

NOTICE: VRF names are case sensitive in both cases: when you create them and when you apply them to layer 3 interfaces, make sure you are using the right capitalization.

Move to interface VLAN X11 and attach it to the VRF **TABLE-X**.

```
Core-1(config)# interface vlan X11
Core-1(config-if-vlan)# vrf attach TABLE-X
Core-1(config-if-vlan)# exit
```

4. Move to interface VLAN X12 and attach it to the VRF **TABLE-X**.

```
Core-1(config)# interface vlan X12
Core-1(config-if-vlan)# vrf attach TABLE-X
Core-1(config-if-vlan)# exit
```

5. Display the Layer 3 interfaces attached to TABLE-X VRF.

```
Core-1(config)# show ip interface brief vrf TABLE-X
Interface           IP Address                Interface Status
                                                link/admin
vlan1111          No Address                    up/up

vlan1112          No Address                    up/up


Core-1(config)#
```

What are the IP addresses of the SVIs?

____

NOTE: When moving a Layer 3 interface (either routing port or SVI) from one VRF to another, it loses all of its IP settings. Therefore, you must configure those parameters again.

6. Assign former IP addresses to interface VLAN X11 and X12.

```
Core-1(config)# interface vlan X11
Core-1(config-if-vlan)# ip address 10.X.11.1/24
Core-1(config-if-vlan)# exit
Core-1(config)# interface vlan X12
Core-1(config-if-vlan)# ip address 10.X.12.1/24
Core-1(config-if-vlan)# exit
Core-1(config)#
```

7. Repeat **step 6**.

```
Core-1(config)# show ip interface brief vrf TABLE-X
Interface         IP Address             Interface Status
                                           link/admin
vlan1111        10.11.11.1/24              up/up

vlan1112        10.11.12.1/24              up/up


Core-1(config)#
```

NOTE: IP connectivity is reestablished in VLANs X11 and X12, however all the typical Layer 3 diagnostic and configuration commands will now be VRF dependent. This means commands will require the VRF name in the command syntax.

8. Display your customer’s routing table. You will need the VRF command extension at the end of the line.

```
Core-1(config)# show ip route vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.11.0/24, vrf TABLE-11
        via  vlan1111,  [0/0],  connected
10.11.11.1/32, vrf TABLE-11
        via  vlan1111,  [0/0],  local
10.11.12.0/24, vrf TABLE-11
        via  vlan1112,  [0/0],  connected
10.11.12.1/32, vrf TABLE-11
        via  vlan1112,  [0/0],  local

Core-1(config)#
```

Are the former IP segments shown in the output?

9. Ping PC-3 and PC-4. You will need the VRF command extension at the end of the line. Ping should be successful.

```
Core-1(config)# do ping 10.X.11.103 vrf TABLE-X
PING 10.11.11.103 (10.11.11.103) 100(128) bytes of data.
108 bytes from 10.11.11.103: icmp_seq=1 ttl=128 time=18.8 ms
108 bytes from 10.11.11.103: icmp_seq=2 ttl=128 time=0.807 ms
108 bytes from 10.11.11.103: icmp_seq=3 ttl=128 time=0.750 ms
108 bytes from 10.11.11.103: icmp_seq=4 ttl=128 time=0.794 ms
108 bytes from 10.11.11.103: icmp_seq=5 ttl=128 time=0.613 ms

--- 10.11.11.103 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4108ms
rtt min/avg/max/mdev = 0.613/4.362/18.847/7.242 ms
Core-1(config)#
```

```
Core-1(config)# do ping 10.X.12.104 vrf TABLE-X
PING 10.11.12.104 (10.11.12.104) 100(128) bytes of data.
108 bytes from 10.11.12.104: icmp_seq=1 ttl=128 time=19.7 ms
108 bytes from 10.11.12.104: icmp_seq=2 ttl=128 time=0.694 ms
108 bytes from 10.11.12.104: icmp_seq=3 ttl=128 time=0.809 ms
108 bytes from 10.11.12.104: icmp_seq=4 ttl=128 time=0.909 ms
108 bytes from 10.11.12.104: icmp_seq=5 ttl=128 time=0.856 ms

--- 10.11.12.104 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4056ms
rtt min/avg/max/mdev = 0.694/4.613/19.797/7.592 ms
```

```
Core-1(config)#
```

TIP: Some diagnostic commands like ping, traceroute, ssh session initiation, etc. are not natively supported in the global configuration context. However, you can import them from manager context by beginning the command with a “do” like in the examples above.

Task 2: Save Your Configurations

Objectives

You will now proceed to save your configuration.

Steps

Core-1 (via PC-1)

1. Save the current Core-1's configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

You have completed Lab 7.2!

AOS-CX Switching Fundamentals

Lab 8: Deploying VRRP

Objectives

Once IP routing was deployed successfully, you approached management and made them aware of how much the network routing relies on Core-1 and how it became a single point of failure in the current infrastructure. You have explained that if Core-1 goes down, VLAN X11 and VLAN X12 will not be able to reach one another. One of them asked you, “how can we fix that?” Your proposal is to deploy a standard First Hop Redundancy Protocol (FHRP) called Virtual Router Redundancy Protocol (VRRP).

Objectives

After completing this lab, you will be able to:

- Enable routing functions on Core-2
- Deploy VRRP on both core switches
- Test failover and failback
- Enable VRRP and MST coordination

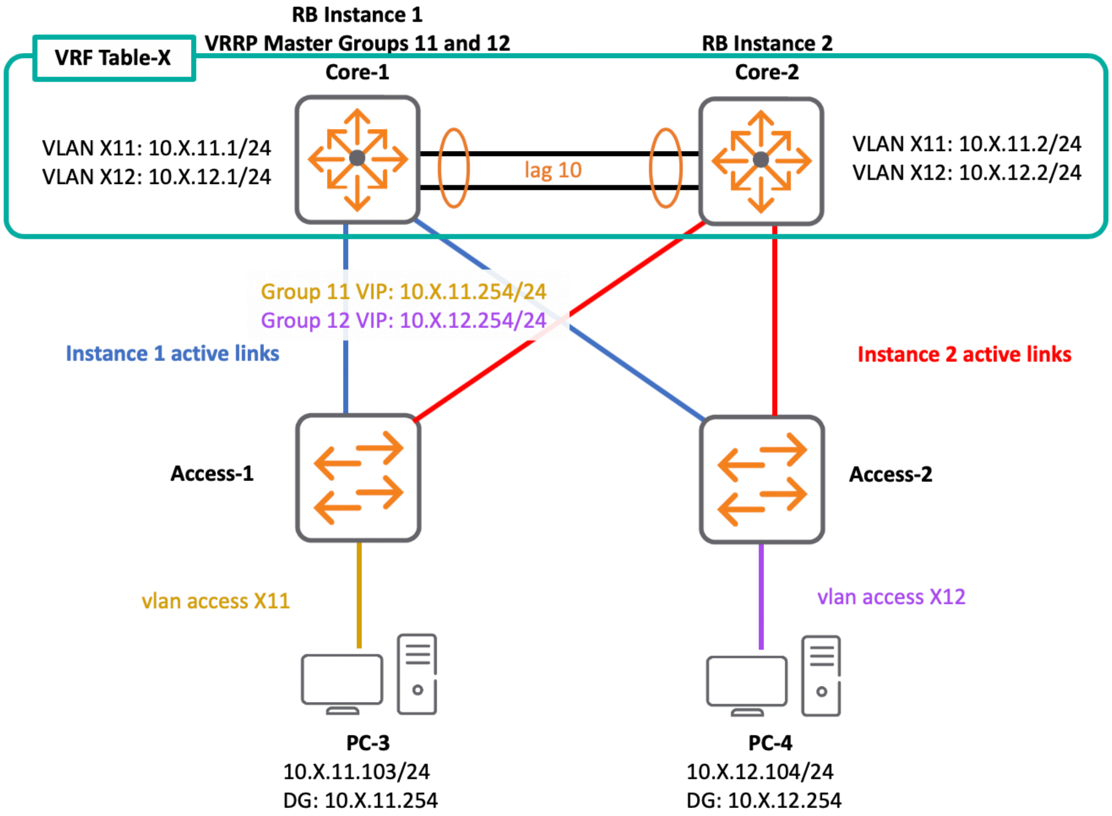

Figure 8-1: Lab Topology

Task 1: Enable IP Settings in Core-2

Objectives

In the following steps you will configure in Core-2 the same VRF and SVIs you already have in Core-1, assign them IP addresses and verify Layer 3 connectivity.

Steps

Core-2 (via PC-1)

1. Open the SSH session to Core-2. Login using **cxfX/aruba123.**

> **NOTE:** Replace the highlighted “X” for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

2. Create **TABLE-X VRF**.

```
Core-2# configure terminal
Core-2(config)# vrf TABLE-X
Core-2(config-vrf)# exit
```

> **NOTICE:** VRF names are case sensitive in both cases: when you create them and when you apply them to layer 3 interfaces, make sure you are using the right capitalization.

3. Create interface **VLAN X11** and attach it to the VRF **TABLE-X**, then assign it IP address **10.X.11.2/24.**

> **NOTE:** Replace the highlighted “X” for your student table number.

```
Core-2(config)# interface vlan X11
Core-2(config-if-vlan)# vrf attach TABLE-X
```

```
Core-2(config-if-vlan)# ip address 10.X.11.2/24
```

4. Create interface **VLAN X12** and attach it to the VRF **TABLE-X**, then assign it IP address **10.X.12.2/24.**

```
Core-2(config)# interface vlan X12
Core-2(config-if-vlan)# vrf attach TABLE-X
Core-2(config-if-vlan)# ip address 10.X.12.2/24
Core-2(config-if-vlan)# end
```

5. Display the Layer 3 interfaces attached to TABLE-X vrf.

```
Core-2# show ip interface vrf TABLE-X

Interface vlan1111 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:3f:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.11.11.2/24
 L3 Counters: Rx Disabled, Tx Disabled

Interface vlan1112 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:3f:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.11.12.2/24
 L3 Counters: Rx Disabled, Tx Disabled
Core-2#
```

What are the IP addresses of the SVIs?

____

6. As a sanity check, confirm you can ping Core-1 using both SVIs.

```
Core-2# ping 10.X.11.1 repetitions 2 vrf TABLE-X
PING 10.11.11.1 (10.11.11.1) 100(128) bytes of data.
108 bytes from 10.11.11.1: icmp_seq=1 ttl=64 time=16.8 ms
108 bytes from 10.11.11.1: icmp_seq=2 ttl=64 time=0.258 ms
```

```
--- 10.11.11.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.258/8.538/16.818/8.280 ms
Core-2#
```

```
Core-2# ping 10.X.12.1 repetitions 2 vrf TABLE-X
PING 10.11.12.1 (10.11.12.1) 100(128) bytes of data.
108 bytes from 10.11.12.1: icmp_seq=1 ttl=64 time=17.9 ms
108 bytes from 10.11.12.1: icmp_seq=2 ttl=64 time=0.286 ms

--- 10.11.12.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.286/9.129/17.973/8.844 ms
Core-2#
```

Task 2: Deploying VRRP

Objectives

Next you will enable a VRRP instance, creating a virtual address and using it as the default gateway on PC-3. You will also track the processes' roles, discover the virtual MAC address used for the Virtual IP, and witness the effect of preemption.

Steps

Core-2 (via PC-1)

1. Open the SSH session to **Core-2**.
2. Move to interface **VLAN X11** and create the VRRP routing process using Group (Virtual Router ID) 11 using the command **vrrp 11 address-family ipv4**.

NOTE: Replace the highlighted "X" for your student table number.

NOTE: VRRP Group number IS NOT table dependent.

```
Core-2# configure terminal
Core-2(config)# interface vlan X11
Core-2(config-if-vlan)# vrrp 11 address-family ipv4
```

3. Define **10.X.11.254** as the virtual IP address then enable the group.

```
Core-2(config-if-vrrp)# address 10.X.11.254 primary
Core-2(config-if-vrrp)# no shutdown
Core-2(config-if-vrrp)# exit
Core-2(config-if-vlan)# exit
```

4. Display the VRRP process information.

```
Core-2(config)# show vrrp interface vlanX11
```

```
VRRP is enabled

Interface vlan1111 - Group 11 - Address-Family IPv4
  State is MASTER
  State duration 01 mins 17.300 secs
  Virtual IP address is 10.11.11.254
  Virtual MAC address is 00:00:5E:00:01:0B
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 100
  Master Router is 10.11.11.2 (local)
  Master Advertisement interval is 1000 msec
  Master Down interval is 3609 msec
Core-2(config)#
```

What is the Virtual Router state?

What is the VIP?

What version is the configuration using?

What is the default priority value?

What is the default advertisement interval?

What is the default master down interval?

IMPORTANT: VRRP needs to be enabled per group and globally in the switch. Since Core switches are a shared resource, the feature has been already enabled using the following command:

```
Core-2(config)# router vrrp enable
```

You will now procced configuring its counterpart Core-1.

Core-1 (via PC-1)

5. Open the SSH session to Core-1.
6. Repeat steps 2 to 3.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11
Core-1(config-if-vlan)# vrrp 11 address-family ipv4
```

```
Core-1(config-if-vrrp)# address 10.X.11.254 primary
Core-1(config-if-vrrp)# no shutdown
```

7. Define priority 254, then enable process globally.

```
Core-1(config-if-vrrp)# priority 254
Core-1(config-if-vrrp)# exit
Core-1(config-if-vlan)# exit
Core-1(config)#
```

8. Display the VRRP process information.

```
Core-1(config)# show vrrp interface vlanX11

VRRP is enabled
```

```
Interface vlan1111 - Group 11 - Address-Family IPv4
  State is MASTER
  State duration 12.501 secs
  Virtual IP address is 10.11.11.254
  Virtual MAC address is 00:00:5E:00:01:0B
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 254
  Master Router is 10.11.11.1 (local)
  Master Advertisement interval is 1000 msec
  Master Down interval is 3007 msec
Core-1(config)#
```

What is the preemption setting and priority value?

____

What is the VRRP router state?

____

ANSWER: Because preemption is enabled, and Core-1's priority is higher than its peer, Core-2, Core-1 became MASTER and Core-2 BACKUP. This means that Core-1 is now the one in charge of advertising the hello packets.

What is the Virtual MAC address?

____

Using your conversion skills acquired on Lab 1, take the last 2 hexadecimal digits of the Virtual MAC and convert them into decimal. What is the result?

____

____

____

Is the result close to any previously defined variable? If so, which one?

____

____

PC-3

9. Move to **PC-3.**
10. Open **Wireshark**, there should be a shortcut on the Desktop.
11. Double click the "**Lab NIC**" entry. That will begin the packet capture on that interface. You will see VRRP packets right away.

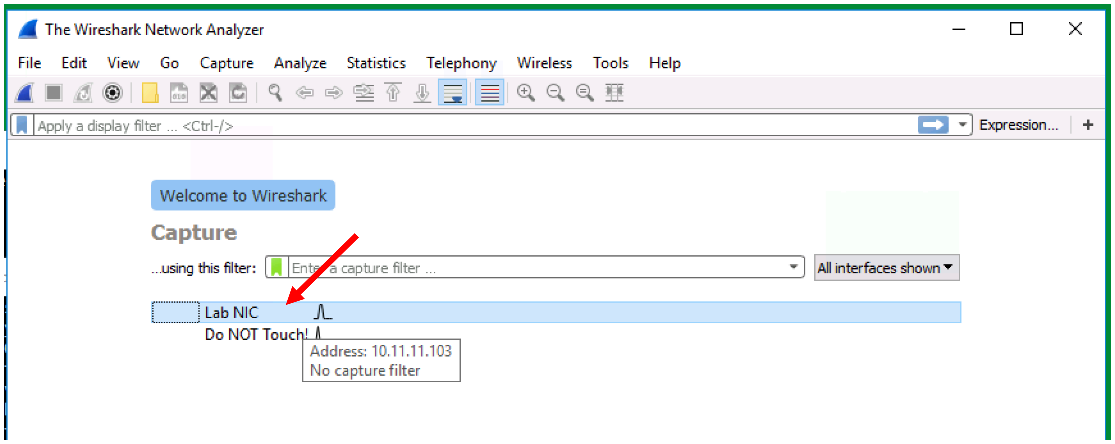

Figure 8-2: Wireshark.

12. Stop the capture.
13. Select any of the VRRP packets.

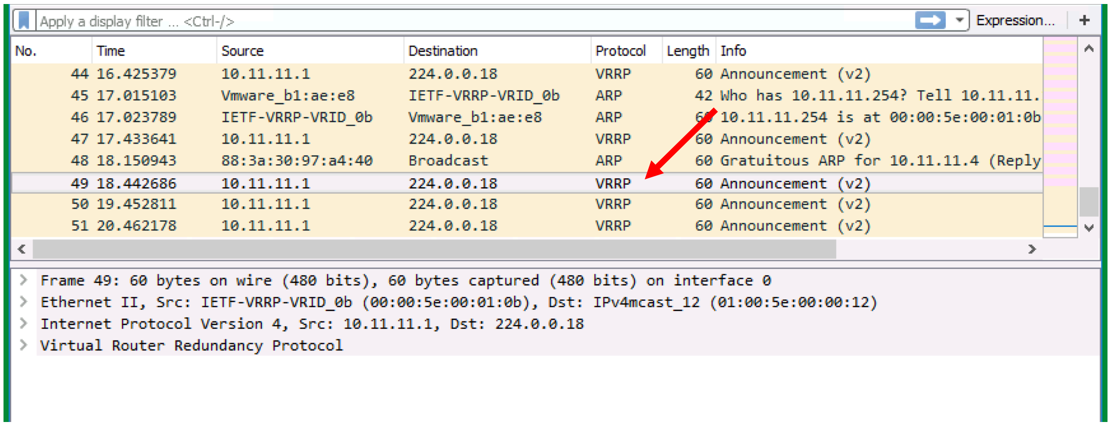

```
No.   Time        Source             Destination        Protocol  Length Info
   44 16.425379   10.11.11.1         224.0.0.18         VRRP          60 Announcement (v2)
   45 17.015103   Vmware_b1:ae:e8    IETF-VRRP-VRID_0b  ARP           42 Who has 10.11.11.254? Tell 10.11.11.
   46 17.023789   IETF-VRRP-VRID_0b  Vmware_b1:ae:e8    ARP           60 10.11.11.254 is at 00:00:5e:00:01:0b
   47 17.433641   10.11.11.1         224.0.0.18         VRRP          60 Announcement (v2)
   48 18.150943   88:3a:30:97:a4:40  Broadcast          ARP           60 Gratuitous ARP for 10.11.11.4 (Reply
   49 18.442686   10.11.11.1         224.0.0.18         VRRP          60 Announcement (v2)
   50 19.452811   10.11.11.1         224.0.0.18         VRRP          60 Announcement (v2)
   51 20.462178   10.11.11.1         224.0.0.18         VRRP          60 Announcement (v2)
```

```
> Frame 49: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
> Ethernet II, Src: IETF-VRRP-VRID_0b (00:00:5e:00:01:0b), Dst: IPv4mcast_12 (01:00:5e:00:00:12)
> Internet Protocol Version 4, Src: 10.11.11.1, Dst: 224.0.0.18
> Virtual Router Redundancy Protocol
```

Figure 8-3: VRRP packet capture

What are source and destination MAC addresses in the Ethernet header?

What kind of address is the destination address?

What are the source and destination addresses in the IP header?

What kind of address is the destination address?

14. Expand the IP header.

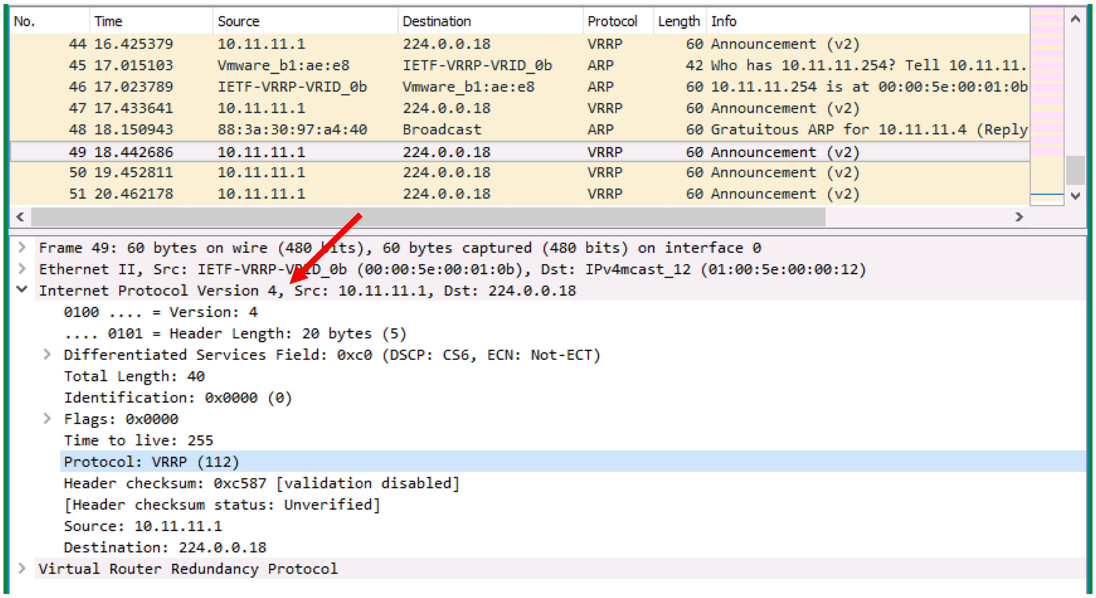

```
No.   Time       Source              Destination          Protocol Length Info
   44 16.425379  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   45 17.015103  Vmware_b1:ae:e8     IETF-VRRP-VRID_0b    ARP       42 Who has 10.11.11.254? Tell 10.11.11.
   46 17.023789  IETF-VRRP-VRID_0b   Vmware_b1:ae:e8      ARP       60 10.11.11.254 is at 00:00:5e:00:01:0b
   47 17.433641  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   48 18.150943  88:3a:30:97:a4:40   Broadcast            ARP       60 Gratuitous ARP for 10.11.11.4 (Reply
   49 18.442686  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   50 19.452811  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   51 20.462178  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
```

```
> Frame 49: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
> Ethernet II, Src: IETF-VRRP-VRID_0b (00:00:5e:00:01:0b), Dst: IPv4mcast_12 (01:00:5e:00:00:12)
v Internet Protocol Version 4, Src: 10.11.11.1, Dst: 224.0.0.18
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0xc0 (DSCP: CS6, ECN: Not-ECT)
    Total Length: 40
    Identification: 0x0000 (0)
  > Flags: 0x0000
    Time to live: 255
    Protocol: VRRP (112)
    Header checksum: 0xc587 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.11.11.1
    Destination: 224.0.0.18
> Virtual Router Redundancy Protocol
```

Figure 8-4: VRRP packet details

What is the IP protocol number?

____

15. Expand the VRRP header.

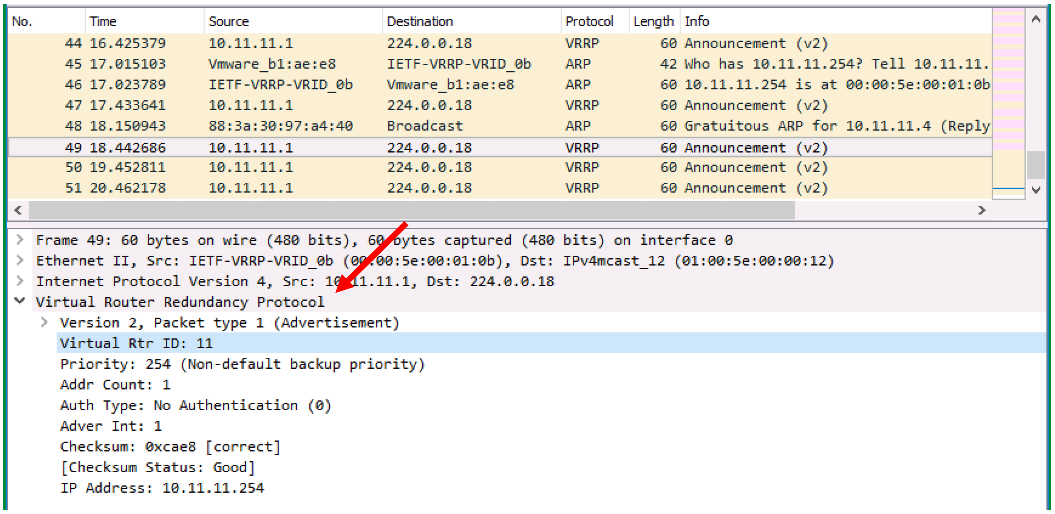

```
No.   Time       Source              Destination          Protocol Length Info
   44 16.425379  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   45 17.015103  Vmware_b1:ae:e8     IETF-VRRP-VRID_0b    ARP       42 Who has 10.11.11.254? Tell 10.11.11.
   46 17.023789  IETF-VRRP-VRID_0b   Vmware_b1:ae:e8      ARP       60 10.11.11.254 is at 00:00:5e:00:01:0b
   47 17.433641  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   48 18.150943  88:3a:30:97:a4:40   Broadcast            ARP       60 Gratuitous ARP for 10.11.11.4 (Reply
   49 18.442686  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   50 19.452811  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
   51 20.462178  10.11.11.1          224.0.0.18           VRRP      60 Announcement (v2)
```

```
> Frame 49: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
> Ethernet II, Src: IETF-VRRP-VRID_0b (00:00:5e:00:01:0b), Dst: IPv4mcast_12 (01:00:5e:00:00:12)
> Internet Protocol Version 4, Src: 10.11.11.1, Dst: 224.0.0.18
v Virtual Router Redundancy Protocol
  > Version 2, Packet type 1 (Advertisement)
    Virtual Rtr ID: 11
    Priority: 254 (Non-default backup priority)
    Addr Count: 1
    Auth Type: No Authentication (0)
    Adver Int: 1
    Checksum: 0xcae8 [correct]
    [Checksum Status: Good]
    IP Address: 10.11.11.254
```

Figure 8-5: VRRP packet details 2

What parameters are familiar to you?

16. Open command prompt and ping the VIP (**10.X.11.254**). Ping should be successful.

NOTE: Replace the highlighted “X” for your student table number.

17. Display the ARP table.

What is the MAC address mapped to the VIP?

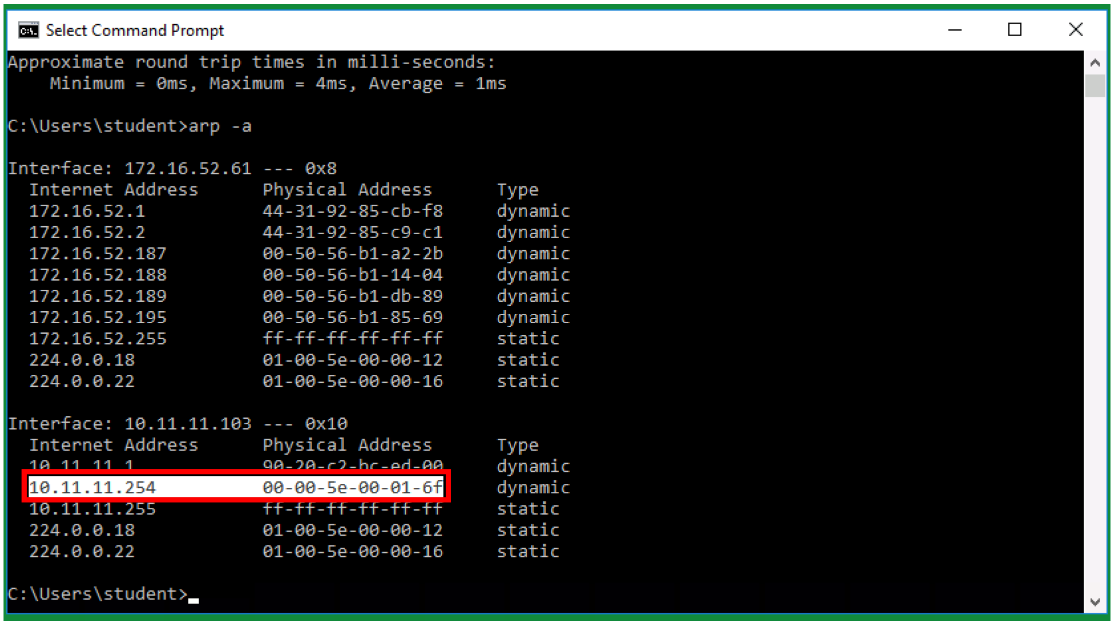

```
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 4ms, Average = 1ms

C:\Users\student>arp -a

Interface: 172.16.52.61 --- 0x8
  Internet Address      Physical Address      Type
  172.16.52.1           44-31-92-85-cb-f8     dynamic
  172.16.52.2           44-31-92-85-c9-c1     dynamic
  172.16.52.187         00-50-56-b1-a2-2b     dynamic
  172.16.52.188         00-50-56-b1-14-04     dynamic
  172.16.52.189         00-50-56-b1-db-89     dynamic
  172.16.52.195         00-50-56-b1-85-69     dynamic
  172.16.52.255         ff-ff-ff-ff-ff-ff     static
  224.0.0.18            01-00-5e-00-00-12     static
  224.0.0.22            01-00-5e-00-00-16     static

Interface: 10.11.11.103 --- 0x10
  Internet Address      Physical Address      Type
  10.11.11.1            90-20-c2-bc-ed-00     dynamic
  10.11.11.254          00-00-5e-00-01-6f     dynamic
  10.11.11.255          ff-ff-ff-ff-ff-ff     static
  224.0.0.18            01-00-5e-00-00-12     static
  224.0.0.22            01-00-5e-00-00-16     static

C:\Users\student>_
```

Figure 8-6: PC-3’s ARP table.

Now that you know how VRRP works, you will proceed with configuring Virtual Router ID 12 for VLAN X12.

Core-1 (via PC-1)

18. Open the SSH session to **Core-1.**
19. **Repeat steps 2 and 3** for **VLAN X12** using **12** and **10.X.12.254** as the VRRP group and VIP respectively.

NOTE: VRRP Group number IS NOT table dependent.

```
Core-1# configure terminal
Core-1(config)# interface vlan X12
Core-1(config-if-vlan)# vrrp 12 address-family ipv4
```

```
Core-1(config-if-vrrp)# address 10.X.12.254 primary
Core-1(config-if-vrrp)# no shutdown
Core-1(config-if-vrrp)# exit
```

Core-2 (via PC-1)

20. Open the SSH session to **Core-2**.
21. Repeat **step 19**.

```
Core-2# configure terminal
Core-2(config)# interface vlan X12
Core-2(config-if-vlan)# vrrp 12 address-family ipv4
```

```
Core-2(config-if-vrrp)# address 10.X.12.254 primary
Core-2(config-if-vrrp)# no shutdown
Core-2(config-if-vrrp)# end
```

22. Run "**show vrrp brief | include vlanX**". Core-2 should be BACKUP of both groups.

```
Core-2# show vrrp brief | include vlanX
 vlan1111      11 IPv4 100 3647      N     Y    BACKUP   10.11.11.1 10.11.11.254
 vlan1112      12 IPv4 100 34        N     Y    BACKUP   10.11.12.1 10.11.12.254
Core-2(config-if-vrrp)#
```

Task 3: Test VRRP Failover

Objectives

In this task you will finally test the resiliency that VRRP can offer to the hosts' default gateway.

Steps

PC-4

1. Access **PC-4**.
2. Change the default gateway in "**Lab NIC**" interface to **10.X.12.254.**

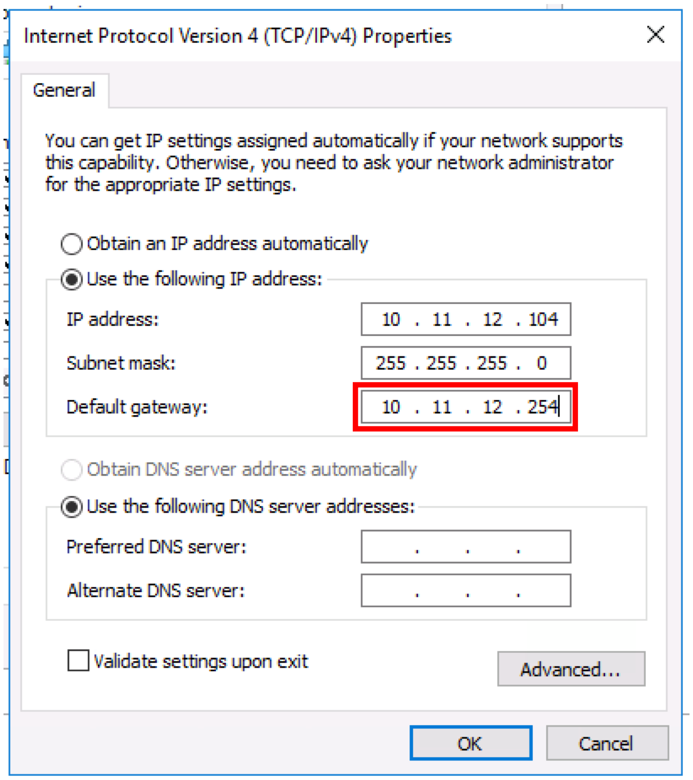

Internet Protocol Version 4 (TCP/IPv4) Properties

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 12 . 104

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 11 . 12 . 254

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 8-7: PC-4's Default gateway.

PC-3

3. Access PC-3.
4. Change the default gateway in **“Lab NIC”** interface to **10.X.11.254**.

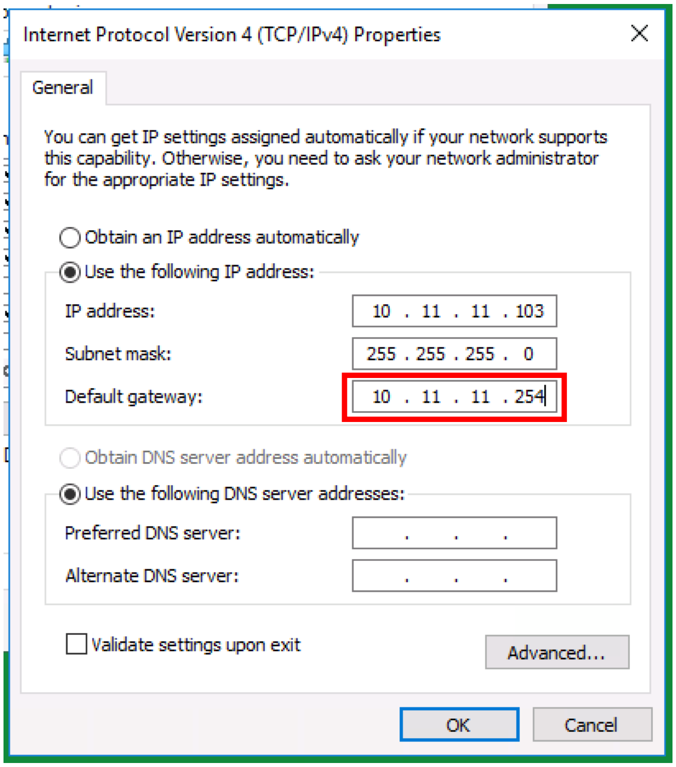

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address: 10 . 11 . 11 . 103

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 11 . 11 . 254

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 8-8: PC-3’s Default gateway.

5. Run a traceroute towards PC-4 (**10.X.12.104**).

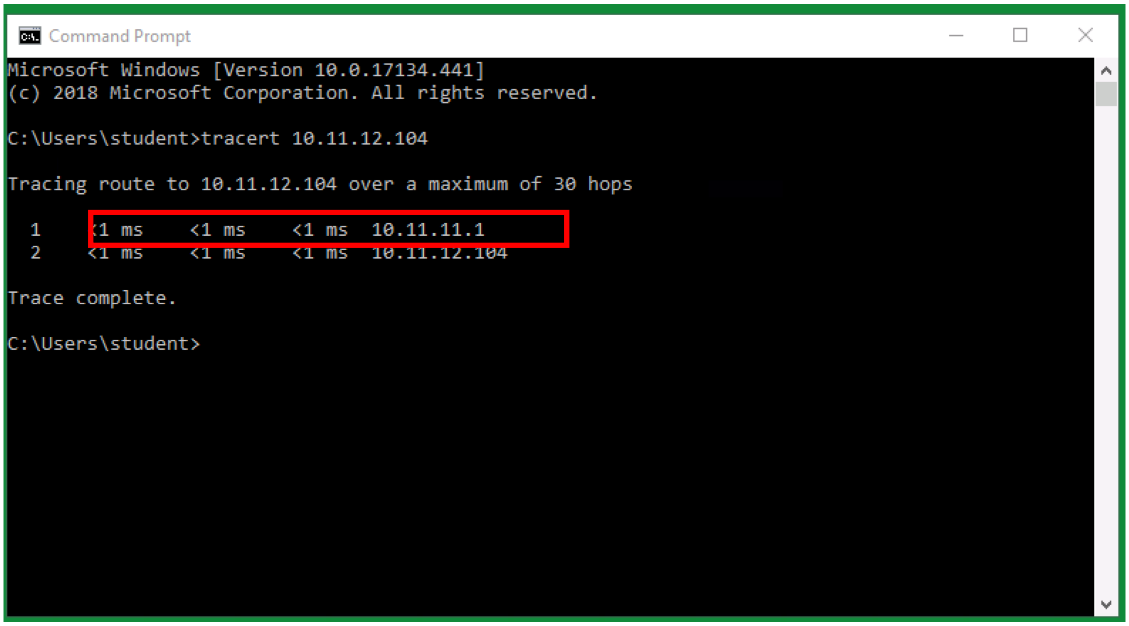

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 10.11.12.104

Tracing route to 10.11.12.104 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.12.104

Trace complete.

C:\Users\student>
```

Figure 8-9: Traceroute 1.

Who is your first hop?

____

NOTE: When an AOS-CX switch receives traceroutes with TTL of 1 with the VRRP MAC address as layer 2 destination, the packet will die as normal (after decreasing TTL by 1), and reply will come from the real IP address of Layer 3 interface the switch received the packet on.

6. Open another command prompt window and run a continuous ping to PC-4 (**10.X.12.104**). Ping should be successful.

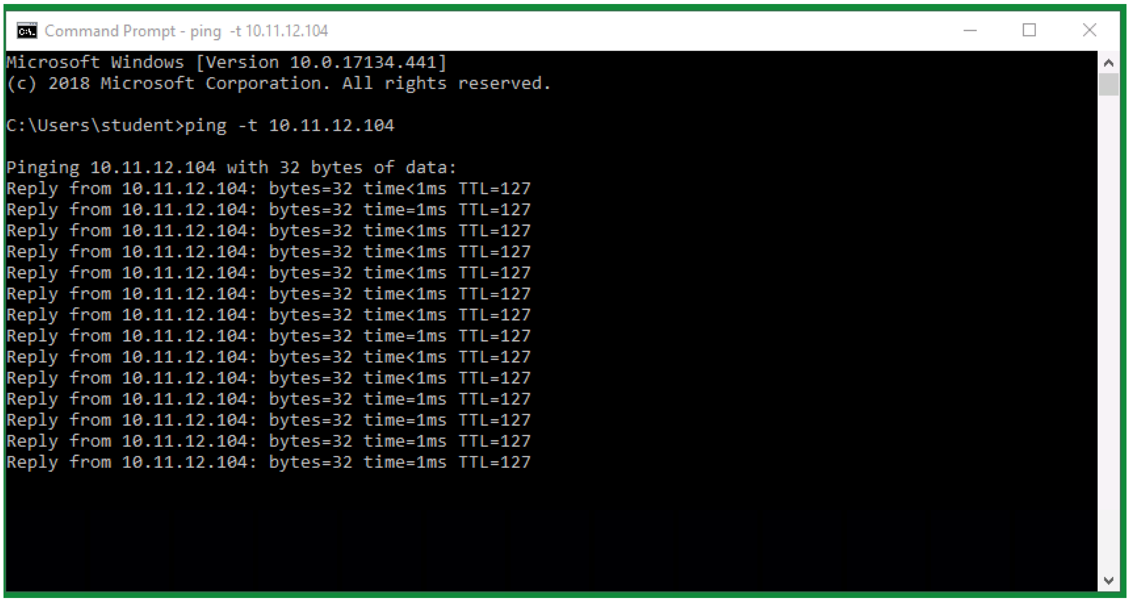

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping -t 10.11.12.104

Pinging 10.11.12.104 with 32 bytes of data:
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
```

Figure 8-10: Ping to PC-4.

Core-1 (via PC-1)

7. Open the SSH session to Core-1.
8. Disable interface VLANs X11 and X12. This will simulate a failure in Core-1 without affecting the other tenants.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11-X12
Core-1(config-if-vlan-<1111-1112>)# shutdown
```

PC-3

9. Move back to PC-3.

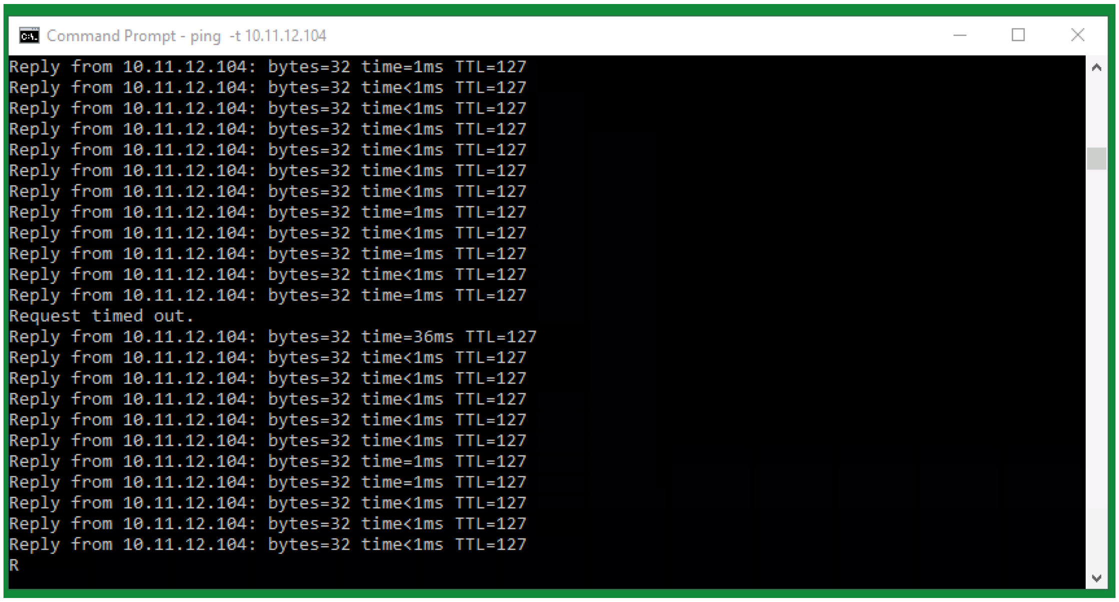

```
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Request timed out.
Reply from 10.11.12.104: bytes=32 time=36ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
R
```

Figure 8-11: Ping to PC-4 2.

How many pings did you missed?

____

Is that what you expected?

____

10. Repeat the traceroute.

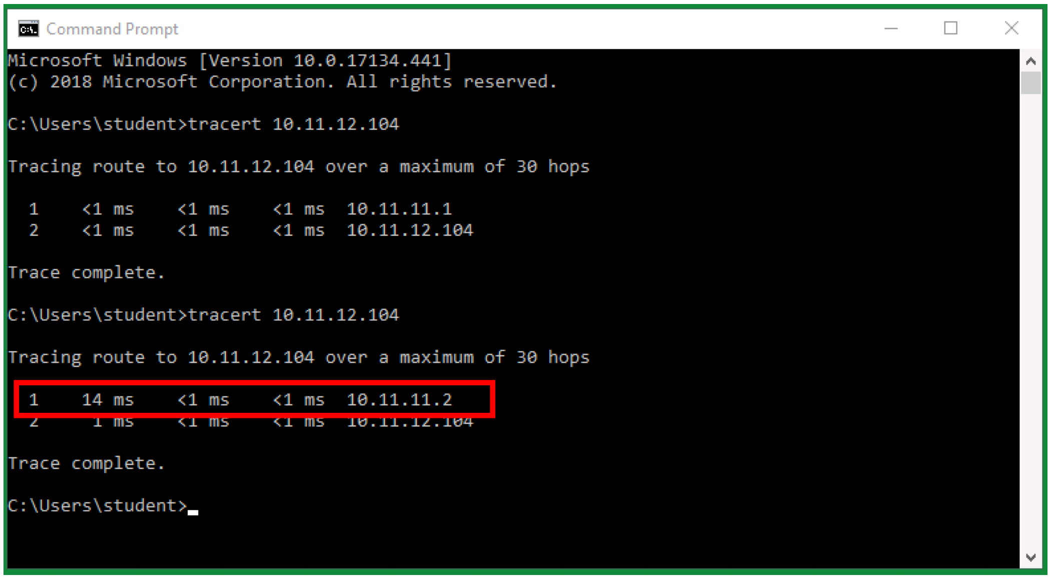

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 10.11.12.104

Tracing route to 10.11.12.104 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.12.104

Trace complete.

C:\Users\student>tracert 10.11.12.104

Tracing route to 10.11.12.104 over a maximum of 30 hops

  1    14 ms    <1 ms    <1 ms  10.11.11.2
  2     1 ms    <1 ms    <1 ms  10.11.12.104

Trace complete.

C:\Users\student>_
```

Figure 8-12: Traceroute 2.

Who is your first hop now?

____

Core-2 (via PC-1)

11. Move back to Core-2.
12. Display the brief version of VRRP. Core-2 should be MASTER on both groups.

```
Core-2# show vrrp brief | include vlanX
 vlan1111     11 IPv4 100 1          N      Y    MASTER  10.11.11.2 10.11.11.254
 vlan1112     12 IPv4 100 2          N      Y    MASTER  10.11.12.2 10.11.12.254
Core-2# show vrrp brief
```

Core-1 (via PC-1)

13. Move back to Core-1
14. Enable interface VLANs X11 and X12.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11-X12
Core-1(config-if-vlan-<1111-1112>)# no shutdown
```

Task 4: VRRP and MST coordination

Objectives

As seen in Task 2, in case of a priority tie the current MASTER remains MASTER, this makes Core-1 control both VIPs under some situations e.g. a power outage when both Core switches go down and Core1 beats Core2 during the boot process.

The problem with this is that Layer 3 load balancing is not guaranteed.

You currently have load sharing at Layer 2 by distributing the different MST instances' root bridges. A best practice is to coordinate both MST and VRRP as seen in figure 8-13. This way, under normal conditions Core-1 is both the root bridge for instance 1 (where VLAN X11 belongs) and the VRRP Master for VLAN X11's VIP. Likewise, Core-2 is both the root bridge for instance 2 (where VLAN X12 belongs) and the VRRP Master for VLAN X12's VIP. The ultimate result is when traffic has to go out the local segment. As soon as traffic hits either Core switch at Layer 2, that device is the gateway in charge of routing the traffic at Layer 3.

The next step raises the priority of Core-2 to achieve the desired behavior.

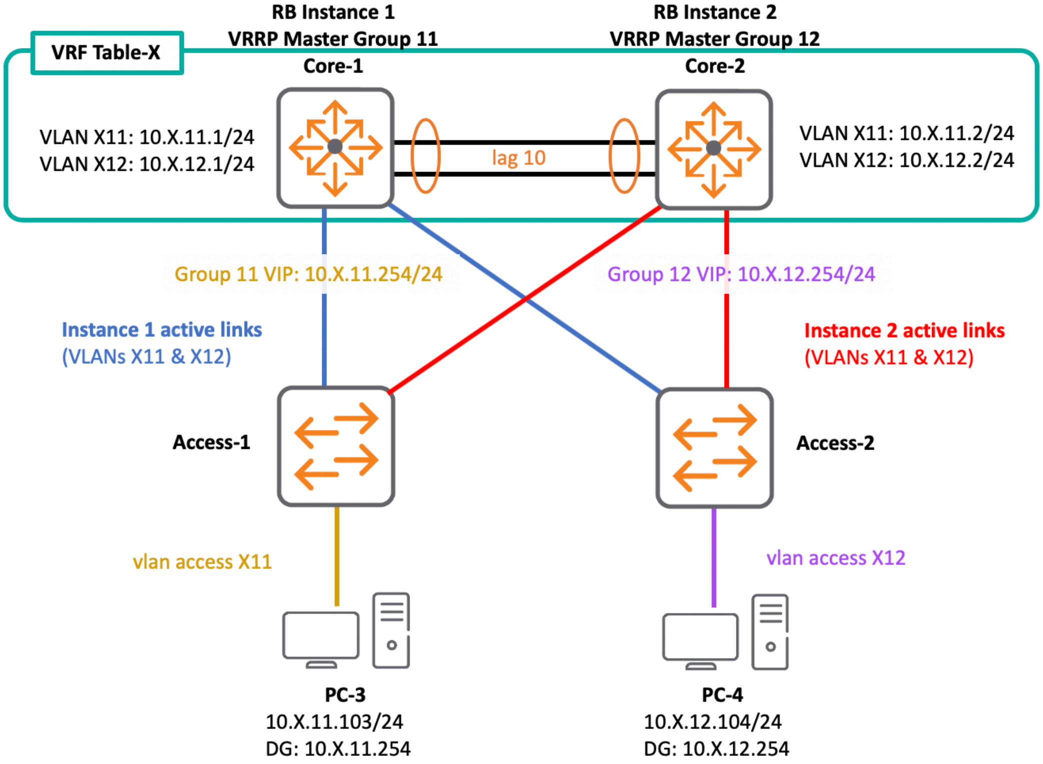

Figure 8-13: MST and VRRP coordination.

Steps

Core-2 (via PC-1)

1. Move back to Core-2.
2. Increase the priority of the VRRP group 12 to 254.

```
Core-2# configure terminal
Core-2(config)# interface vlan X12
Core-2(config-if-vlan)# vrrp 12 address-family ipv4
Core-2(config-if-vrrp)# priority 254
Core-2(config-if-vrrp)# end
```

3. Display the VRRP process information. Core-2 should be BACKUP of group 11 and MASTER of 12.

```
Core-2# show vrrp brief | include vlanX
vlan1111     11 IPv4 100 1659      N     Y   BACKUP  10.11.11.1 10.11.11.254
vlan1112     12 IPv4 254 1667      N     Y   MASTER  10.11.12.2 10.11.12.254
Core-2#
```

Task 5: Save Your Configurations

Objectives

You will now proceed to save your configuration.

Steps

Core-1 and Core-2 (via PC-1)

4. Save the current Core-1 and Core-2 configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

You have completed Lab 8!

AOS-CX Switching Fundamentals

Lab 9: Subnetting and VLSM

Overview

BigStartup has plans to expand the network starting with the acquisition of Internet links from two different carriers, followed by adding a Server Switch, and investing in an Aruba Instant Solution.

You have been asked to interconnect the Core Switches with a Perimeter firewall pair that will connect to these ISP links; using non-/24 prefixes. They also want you to reserve two IP segments for connections to the Server Switch and another one for hosting up to 500 WiFi clients. Therefore, you have decided to review and practice subnetting before jumping into any configuration.

Objectives

After completing this lab, you will be able to:

- Subnet Class A, B and C networks into smaller IP segments
- Calculate the total number of networks and hosts that a subnet process generates
- Identify the Network ID and broadcast IP of a Subnet
- Identify assignable IP address in a particular Subnet

Task 1: Class A Subnetting

Objectives

Subnet the prefix using the information below:

Network Address: **43.0.0.0**
Number of needed Subnets: **9**

Steps

1. List all subnets in table 9-1 down below.

What is the address class?

____

What is the default subnet mask?

____

What is the required subnet mask?

____

How many subnets will be generated with equal length subnet mask?

____

What is the total number of assignable addresses per subnet?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-1: Subnetting Task 1

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

Task 2: Class B Subnetting

Objectives

Subnet the prefix using the information below:

IP Address: **132.89.5.10**
Number of needed Subnets: **20**

Steps

1. List all subnets in table 9-2 down below

What network does the address belong to?

____

What is the address class?

____

What is the default subnet mask?

____

What is the required subnet mask?

____

How many subnets will be generated with equal length subnet mask?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-2: Subnetting Task 2

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

Task 3a: Class C Subnetting Part 1

Objectives

Subnet the prefix using the information below:

Network Address: **192.168.1.0**
Number of needed assignable host addresses: **2**

Steps

1. List the first 4 subnets and the last one in table 9-3

What is the address class?

____

What is the default subnet mask?

____

What is the required subnet mask?

____

How many subnets will be generated with equal length subnet mask?

____

What is the total number of assignable addresses per subnet?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-3: Subnetting Task 3a

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

Task 3b: Class C Subnetting Part 2 (optional)

Objectives

Subnet the prefix using the information below:

Network Address: **199.209.0.0**
Number of needed Subnets: **7**

Steps

1. List all subnets in table 9-4 down below.

What is the address class?

____

What is the default subnet mask?

____

What is the required subnet mask?

____

How many subnets will be generated with equal length subnet mask?

____

What is the total number of assignable addresses per subnet?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-4: Subnetting Task 3b

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

Task 4a: VLSM Prefixes

Objectives

Subnet the prefix using the information below:

Network Address: **10.0.0.0**
Number of needed assignable host addresses: **254**

Steps

1. List 1st, 2nd, 3rd, 21st , 22nd, and the 101st subnets in table 9-5 down below

What is the address class?

____

What is the default subnet mask?

____

What is the required subnet mask?

____

How many subnets will be generated with equal length subnet mask?

____

What is the total number of assignable addresses per subnet?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-5: Subnetting exercise 4a

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

Task 4b: VLSM – Point to Point Segments

Objectives

Subnet the prefix using the information below:

Take the first /24 subnet of exercise 4a and subnet it again with segments that support up to 2 assignable addresses.

Steps

1. List the first 5 subnets in table 9-6.

   What is the required subnet mask?

____

   How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-6: Subnetting exercise 4b

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

IMPORTANT: It is always a best practice to deploy a /30 prefix when the segment will be used on a link (physical or virtual) that only interconnects two Layer 3 devices e.g. Ethernet links between two routers or multilayer switches, GRE tunnels, serial links, etc.

Task 4c: VLSM – Grouping Two Subnets (optional)

Objectives

Subnet the prefix using the information below:

Combine subnets 21 and 22 of exercise 4a into a single one that supports 500 hosts.

Steps

1. List the resulting subnet in table 9-6.

   What is the total number of assignable addresses?

____

   What is the required subnet mask?

____

   How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-6: Subnetting Task 4c

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

IMPORTANT: Multiple contiguous subnets can be combined together into a larger one in order to provide more assignable addresses within the same segment or, for summarization purposes when using static routes or dynamic routing protocols.

Task 4d: VLSM – Loopback Segments (optional)

Objectives

Subnet the prefix using the information below:

Use the 101st subnet of exercise 4a and subnet it again with segments that support up to 1 host address.

Steps

1. List the first 5 subnets in table 9-7.

What is the required subnet mask?

____

How many subnets will be generated with this new subnet mask?

____

How many bits were borrowed from the host portion in the default mask for creating subnets?

____

Table 9-7: Subnetting exercise 5d

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address

IMPORTANT: /32 prefixes with a single host address can be used on loopback interfaces for connectivity tests, management and routing protocols for fine tuning (OSPF Router ID reachability, iBGP and eBGP multihop peering, etc). The device that owns the address is the only one with direct access to it unless you tell other devices how to reach it with static or dynamic routing since the address will be it's segment itself!

It is always a good practice to reserve a range of addresses of your IP address scheme for this purpose and allocate one of them to each Layer 3 device in the network.

You have completed Lab 9!

AOS-CX Switching Fundamentals

Lab 10: Static Routes

Overview

The goal of the following tasks is to complete the dual-homed Internet Service deployment for BigStartup. The customer wants load balancing across both carriers and redundancy in case of failure. They want assurance that if either link fails, traffic can still go out through the alternate ISP. This will require the configuration of static and floating routes, which you will apply on the Core switches.

Objectives

After completing this lab, you will be able to:

- Configure Core switches to Perimeter Firewall links using a /30 prefix
- Calculate and deploy Variable Length Subnet Mask (VLSM) prefixes
- Configure static routes
- Add a default route into the routing table for providing internet access
- Manipulate administrative distances in order to configure floating routes
- Validate proper load sharing and failover

NOTE: IP prefix is an aggregation of IP addresses and is usually used to refer to an IP network or subnet in general.

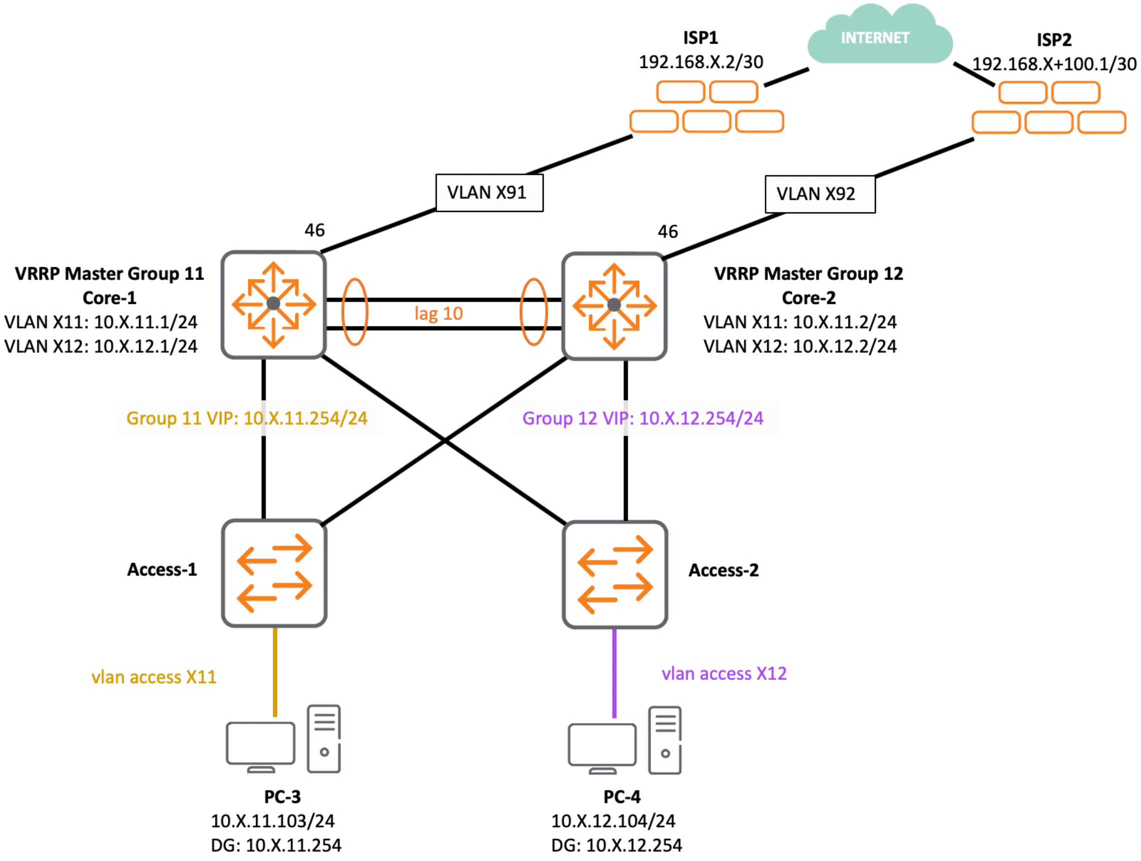

Figure 10-1: Lab Topology

Task 1: Add Links to ISPs

Objectives

In this task, you will prepare the network for future changes such as the addition of internet connections by assigning the /30 segments you calculated on Lab 9.1 Task 3b to VLANs X91 and X92 on Core-1 and Core-2 respectively.

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1. Login using **cxfX/aruba123.**

> **NOTE:** Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

2. Create **VLAN X91** and add the name **TX-ISP-1**.

> **NOTE:** Replace the highlighted "X" for your student table number.

```
Core-1# configure terminal
Core-1(config)# vlan X91
Core-1(config-vlan-1191)# name TX-ISP-1
Core-1(config-vlan-1191)# exit
```

3. Create interface **VLAN X91** and map it to vrf **TABLE-X**.

```
Core-1(config)# interface vlan X91
Core-1(config-if-vlan)# vrf attach TABLE-X
Core-1(config-if-vlan)#
```

4. Assign IP address **192.168.X.1/30**.

```
Core-1(config-if-vlan)# ip address 192.168.X.1/30
Core-1(config-if-vlan)# exit
```

5. Move to port 1/1/46 and allow **VLAN X91**.

```
Core-1(config)# interface 1/1/46
Core-1(config-if)# vlan trunk allowed X91
Core-1(config-if)# exit
```

6. Confirm you can ping ISP1 (**192.168.X.2**).

```
Core-1(config)#do ping 192.168.X.2 vrf TABLE-X
PING 192.168.11.2 (192.168.11.2) 100(128) bytes of data.
108 bytes from 192.168.11.2: icmp_seq=1 ttl=64 time=13.4 ms
108 bytes from 192.168.11.2: icmp_seq=2 ttl=64 time=0.172 ms
108 bytes from 192.168.11.2: icmp_seq=3 ttl=64 time=0.162 ms
108 bytes from 192.168.11.2: icmp_seq=4 ttl=64 time=0.159 ms
108 bytes from 192.168.11.2: icmp_seq=5 ttl=64 time=0.154 ms

--- 192.168.11.2 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4049ms
rtt min/avg/max/mdev = 0.154/2.810/13.405/5.297 ms
Core-1(config-if)# end
```

TIP: Some commands like copy, ping, or traceroute are not natively available at configuration context, however you can use the “do” command in order to import them from privileged context.

Core-2 (via PC-1)

7. Open the SSH session to Core-2
8. Repeat steps 2 to 6 using VLAN X92, TX_ISP-2 as description and **192.168.X+100.2/30** as the IP address.

NOTE: Replace the highlighted “X” for your student table number, e.g. “vlan 492” and “192.168.104.2/30” for table 4, or “vlan 1192” and “192.168.111.2/30” for table 11 like in the example below.

```
Core-2# configure terminal
Core-2(config)# vlan X92
Core-2(config-vlan-1192)# name TX_ISP-2
Core-2(config-vlan-1192)# exit
```

```
Core-2(config)# interface vlan X92
Core-2(config-if-vlan)# vrf attach TABLE-X
Core-2(config-if-vlan)#
```

```
Core-2(config-if-vlan)# ip address 192.168.X+100.2/30
Core-2(config-if-vlan)# exit
```

```
Core-2(config)# interface 1/1/46
Core-2(config-if)# vlan trunk allowed X92
Core-2(config-if)# exit
```

```
Core-2(config)# do ping 192.168.X+100.1 vrf TABLE-X
PING 192.168.111.1 (192.168.111.1) 100(128) bytes of data.
108 bytes from 192.168.111.1: icmp_seq=1 ttl=64 time=11.6 ms
108 bytes from 192.168.111.1: icmp_seq=2 ttl=64 time=0.216 ms
108 bytes from 192.168.111.1: icmp_seq=3 ttl=64 time=0.229 ms
108 bytes from 192.168.111.1: icmp_seq=4 ttl=64 time=0.180 ms
108 bytes from 192.168.111.1: icmp_seq=5 ttl=64 time=0.178 ms

--- 192.168.111.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4108ms
rtt min/avg/max/mdev = 0.178/2.488/11.640/4.576 ms
Core-2(config)# end
```

Task 2: Adding Static Routes

Objectives

Right now, the links between the Core Switches and Perimeter Firewalls are up and running, however internet access is not available yet. In this task you will add static routes in order to send all non-local traffic to the carriers who will take care of the delivery process. Core-1 will be pointing to ISP1 and Core-2 will point to ISP2 in order to achieve a load balancing effect.

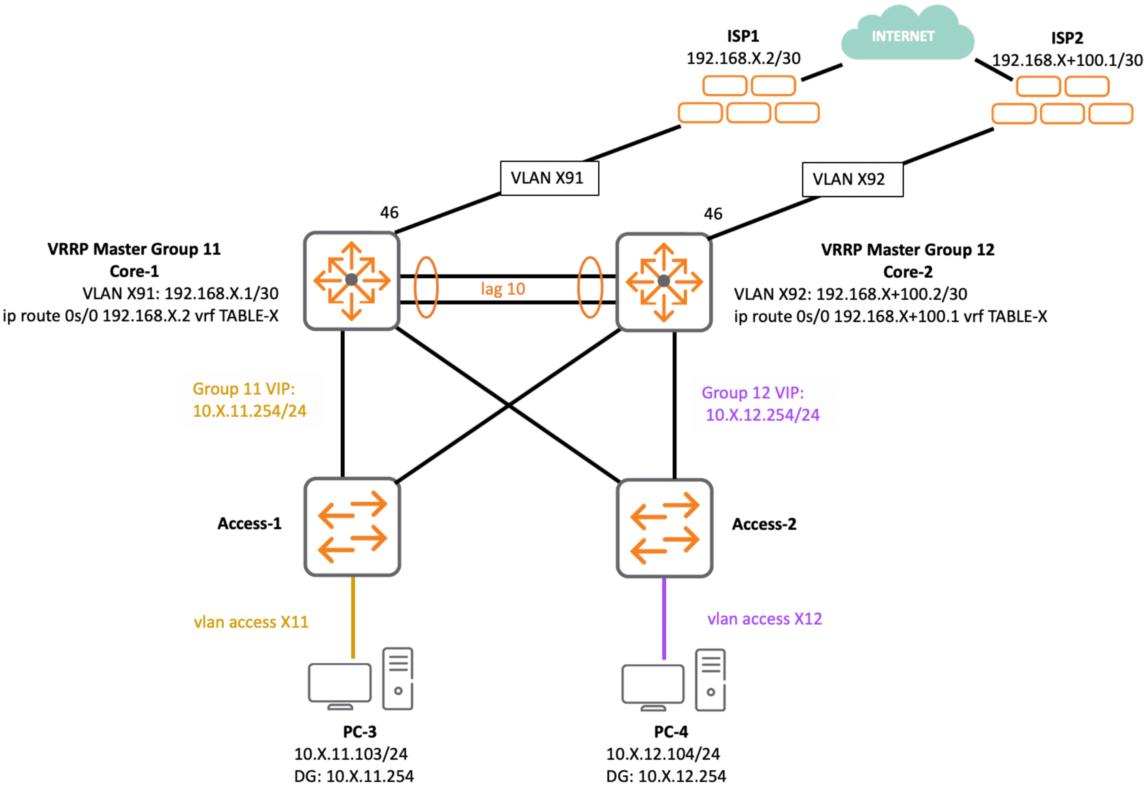

Figure 10-2: Task 2 topology

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1.
2. Create a static default route (also known as 0’s prefix) pointing to ISP-1 (**192.168.X.2**) on **TABLE-X**.

```
Core-1# configure terminal
Core-1(config)# ip route 0.0.0.0/0 192.168.X.2 vrf TABLE-X
Core-1(config)#
```

3. Use **show ip route static vrf** and validate the route is listed.

```
Core-1(config)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  192.168.11.2,  [1/0],  static

Core-1(config)#
```

What is the metric value and what is it for?

____

____

What is the distance value and what is it for?

____

____

4. Ping the 8.8.8.8 IP address. Ping should be successful.

```
Core-1(config)# do ping 8.8.8.8 datagram-size 32 vrf TABLE-X
PING 8.8.8.8 (8.8.8.8) 32(60) bytes of data.
40 bytes from 8.8.8.8: icmp_seq=1 ttl=51 time=8.09 ms
40 bytes from 8.8.8.8: icmp_seq=2 ttl=51 time=7.84 ms
40 bytes from 8.8.8.8: icmp_seq=3 ttl=51 time=7.82 ms
40 bytes from 8.8.8.8: icmp_seq=4 ttl=51 time=8.00 ms
40 bytes from 8.8.8.8: icmp_seq=5 ttl=51 time=8.07 ms

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 7.822/7.968/8.097/0.160 ms
Core-1(config)#
```

TIP: In addition to specifying the VRF, outbound ICMP echo packets can be manipulated by using the **ping** command followed by:

```
   data-fill      Specify the ping packet data pattern in hexadecimal  digits
   datagram-size  Specify the size of ping datagram
   interval       Specify the interval between successive ping requests
   ip-option      Specify an IP option to be used in ping packet
   repetitions    Specify the number of ping packets to be sent
   source         Specify the ping source IPv4 address or source interface
   timeout        Specify the ping timeout in seconds
   tos            Specify IP TOS to be used in ping request
   vrf            Specify the VRF to use
```

If there is no prefix in the routing table for the 8.8.8.8 IP address. What prefix is taking care of routing this traffic?

____

PC-3

5. Access **PC-3** and open a command prompt.
6. Ping the **8.8.8.8** IP address.

Is ping successful?

____

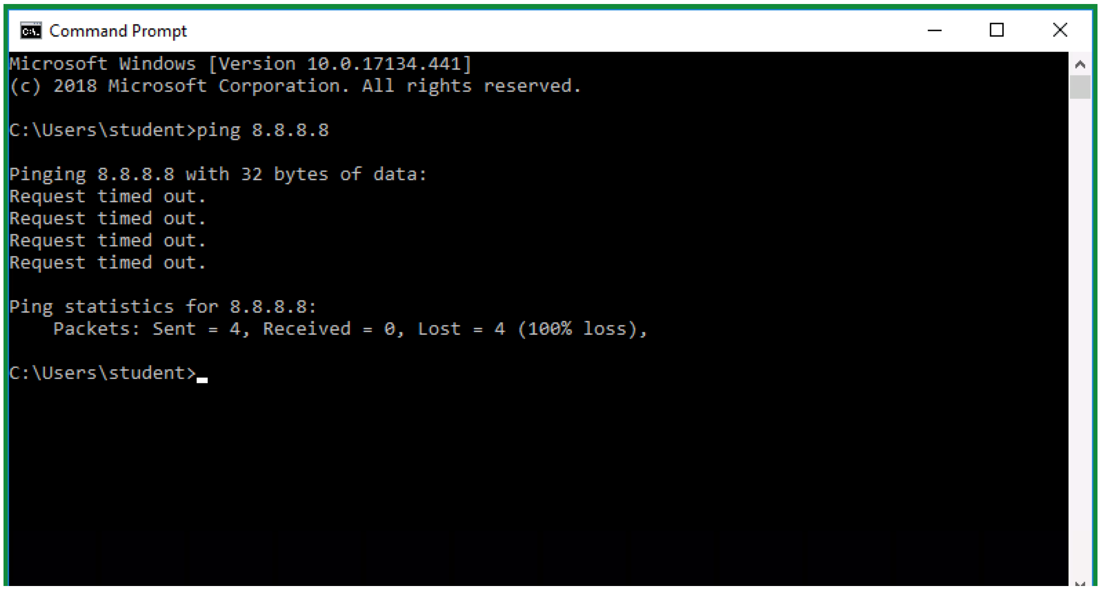

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 8.8.8.8

Pinging 8.8.8.8 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 8.8.8.8:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>
```

Figure 10-3: Failed ping

7. Attempt a traceroute to the same address.

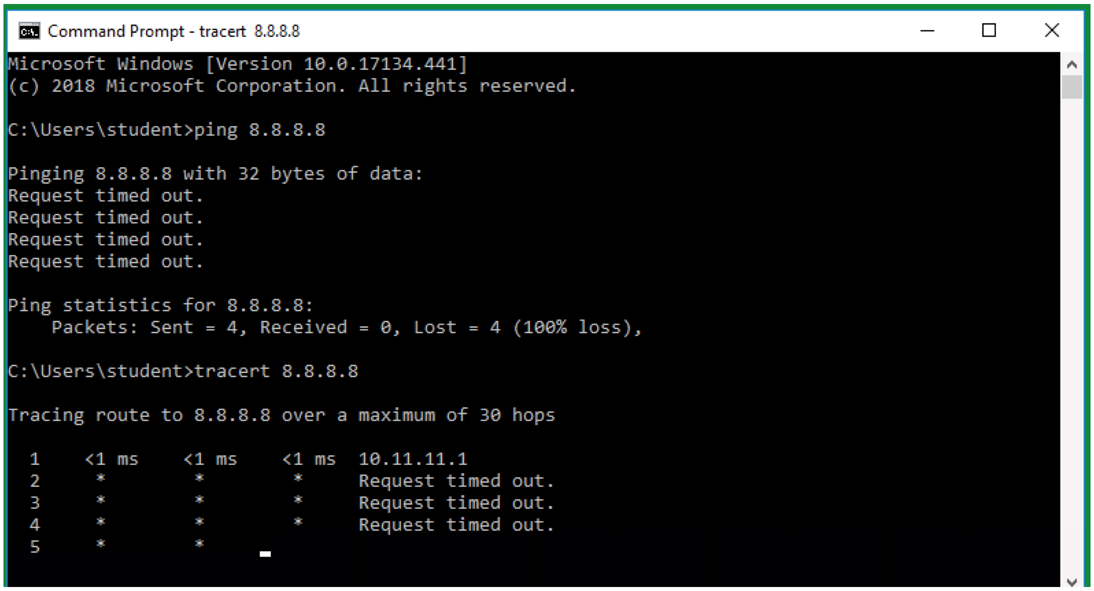

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 8.8.8.8

Pinging 8.8.8.8 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 8.8.8.8:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2     *        *        *     Request timed out.
  3     *        *        *     Request timed out.
  4     *        *        *     Request timed out.
  5     *        *
```

Figure 10-4: Traceroute

What is the last hop your trace is reaching?

What is the last hop your trace is reaching?

NOTE: There could be many reasons why the ping is not working:

a) An ACL in the firewall that filters the packets out.
b) The lack of Network Address Translation (NAT) which sends the packets with the original source IP address making it impossible for the destination to properly respond back to you from the internet.
c) A missing route for your local segment (10.1.0.0/16) in the service provider equipment causing it to drop the returning traffic or route it somewhere else.

At this point any of them is possible, however since Core-1 was able to reach the 8.8.8.8 address then it is most likely that the ISP device (Perimeter Firewall) does not contain your prefix in its routing table. After all, you must remember that when testing access from Core-1, packets had the 192.168.X.1 source IP address which is a segment ISP1 implicitly knows (connected network). On the other hand, packets sent by PC-3 had the 10.X.11.103 address, therefore, you must make sure the carrier has this route in their device pointing to Core-1's IP address as the next hop in VLAN X91.

You have contacted ISP1 and asked if their device was setup properly, ensuring at a minimum the 10.X.11.0/24 and 10.X.12.0 were included in its routing table. After validating the request the ISP realizes that the on-site device is using its own 0's prefix to forward traffic to those segments.

```
ISP1-ServerSwitch# show ip route 10.X.11.0 vrf CXF_ISP1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf CXF_ISP1
        via  192.168.253.254,  [1/0],  static

ISP1-ServerSwitch# show ip route 10.X.12.0 vrf CXF_ISP1

Displaying ipv4 routes selected for forwarding
```

```
'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf CXF_ISP1
        via  192.168.253.254,  [1/0],  static

ISP1-ServerSwitch#
```

To solve this, you request the ISP to add the network 10.X.0.0/16 pointing to 192.168.X.1 IP address (Core-1) as the next hop.

NOTE: In the next steps you will pretend to be the ISP1 technician.

<——————— Begins ISP1 configuration ———————>

ISP1 (via PC-1)

8. Using Putty, open an SSH session to ISP1.

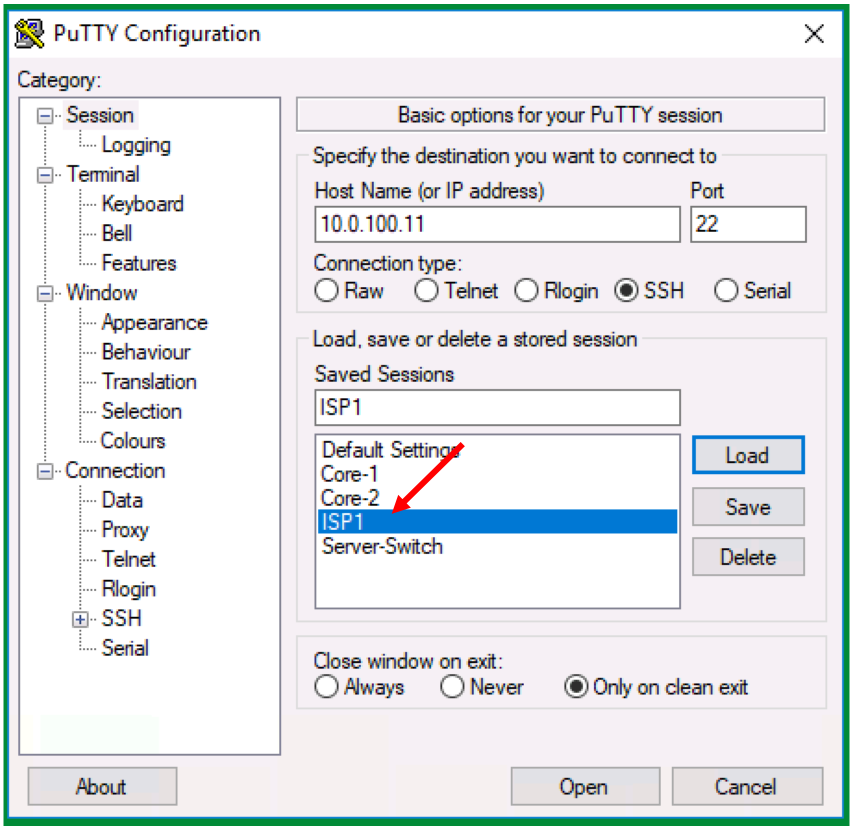

PuTTY Configuration

Basic options for your PuTTY session

Specify the destination you want to connect to

Host Name (or IP address): 10.0.100.11

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Load, save or delete a stored session

Saved Sessions: ISP1

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 10-5: SSH to ISP1

9. Login using username: **cxfX/aruba123.**

> **NOTE:** Replace the highlighted “X” with your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

10. Configure the missing static route: **10.X.0.0/16** via **192.168.X.1** on ISP1 VRF.

```
ISP1-ServerSwitch# configure terminal
ISP1-ServerSwitch(config)# ip route 10.X.0.0/16 192.168.X.1 vrf CXF_ISP1
ISP1-ServerSwitch(config)# end
```

11. Use the **show ip route** command for validating there is an entry in the routing table for properly forwarding traffic to 10.X.11.0/24 and 10.X.12.0/24.

```
ISP1-ServerSwitch# show ip route 10.X.11.0 vrf CXF_ISP1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.0.0/16, vrf CXF_ISP1
        via  192.168.11.1,  [1/0],  static

ISP1-ServerSwitch# show ip route 10.X.12.0 vrf CXF_ISP1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.0.0/16, vrf CXF_ISP1
        via  192.168.11.1,  [1/0],  static

ISP1-ServerSwitch#
```

Is Core-1’s IP address the next-hop in the entries?

12. Close the putty session.

<——————————————— Ends ISP1 configuration ———————————————>

PC-3

13. Move back to PC-3.
14. Ping the **8.8.8.8** IP address, then run a traceroute.

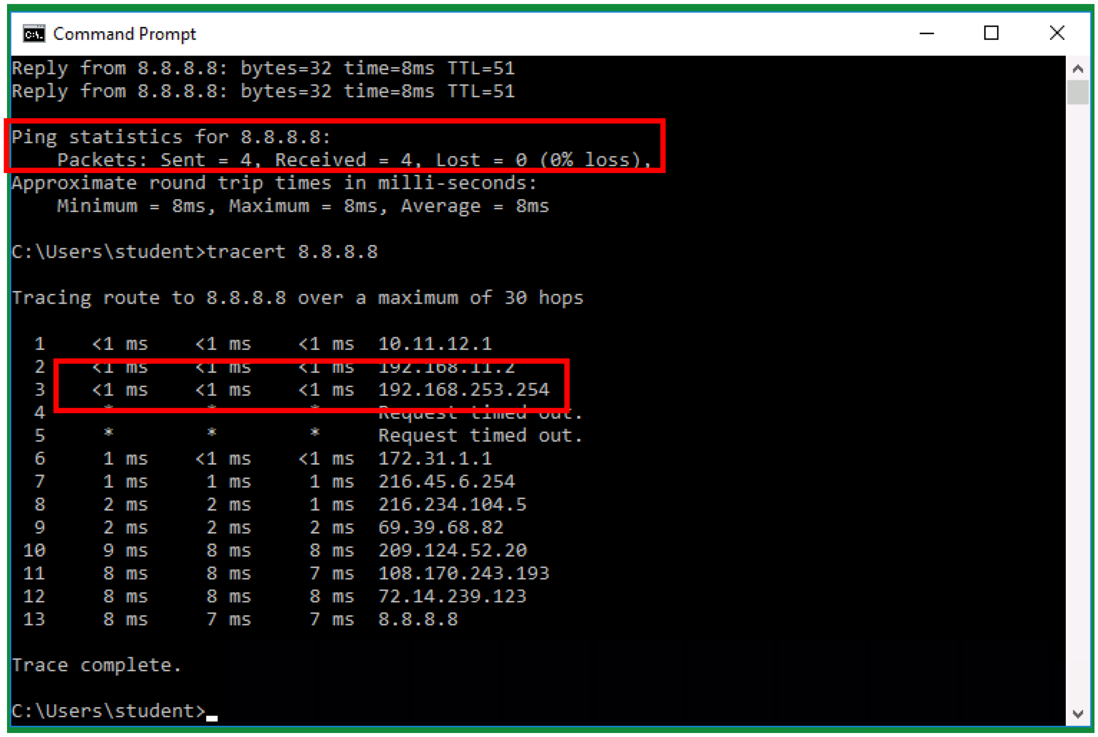

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51

Ping statistics for 8.8.8.8:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 8ms, Maximum = 8ms, Average = 8ms

C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.12.1
  2    <1 ms    <1 ms    <1 ms  192.168.11.2
  3    <1 ms    <1 ms    <1 ms  192.168.253.254
  4     *        *        *     Request timed out.
  5     *        *        *     Request timed out.
  6     1 ms    <1 ms    <1 ms  172.31.1.1
  7     1 ms     1 ms     1 ms  216.45.6.254
  8     2 ms     2 ms     1 ms  216.234.104.5
  9     2 ms     2 ms     2 ms  69.39.68.82
 10     9 ms     8 ms     8 ms  209.124.52.20
 11     8 ms     8 ms     7 ms  108.170.243.193
 12     8 ms     8 ms     8 ms  72.14.239.123
 13     8 ms     7 ms     7 ms  8.8.8.8

Trace complete.

C:\Users\student>_
```

Figure 10-6: Ping Successful

Is ping successful?

____

IMPORTANT: In IP networking, most communications are bidirectional, therefore adding a route with the destination prefix on the layer 3 device next to the source, is just as important as adding a route with the source prefix on the device next to the destination.

If NAT isn't used, then all layer 3 devices in between the source and the destination must have both prefixes in their routing tables as well.

What are the first and second hops?

Core-2 (via PC-1)

15. Open the SSH session to Core-2.
16. Repeat steps 2 to 4 using **192.168.X+100.1** as your next hop.

```
Core-2# configure terminal
Core-2(config)# ip route 0.0.0.0/0 192.168.X+100.1 vrf TABLE-X
Core-2(config)#
```

```
Core-2(config)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  192.168.111.1,  [1/0],  static

Core-1(config)#
```

What is the next hop?

What would happen if that device goes down?

```
Core-2(config)# do ping 8.8.8.8 datagram-size 32 vrf TABLE-X
PING 8.8.8.8 (8.8.8.8) 32(60) bytes of data.
40 bytes from 8.8.8.8: icmp_seq=1 ttl=52 time=8.45 ms
40 bytes from 8.8.8.8: icmp_seq=2 ttl=52 time=8.14 ms
40 bytes from 8.8.8.8: icmp_seq=3 ttl=52 time=7.88 ms
40 bytes from 8.8.8.8: icmp_seq=4 ttl=52 time=8.01 ms
```

```
40 bytes from 8.8.8.8: icmp_seq=5 ttl=52 time=7.91 ms

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 7.889/8.082/8.459/0.216 ms
Core-2(config)#
```

PC-4

17. Access PC-4 and open a command prompt.
18. Ping the **8.8.8.8** IP address. Ping should be successful
19. Run a traceroute to **8.8.8.8**.

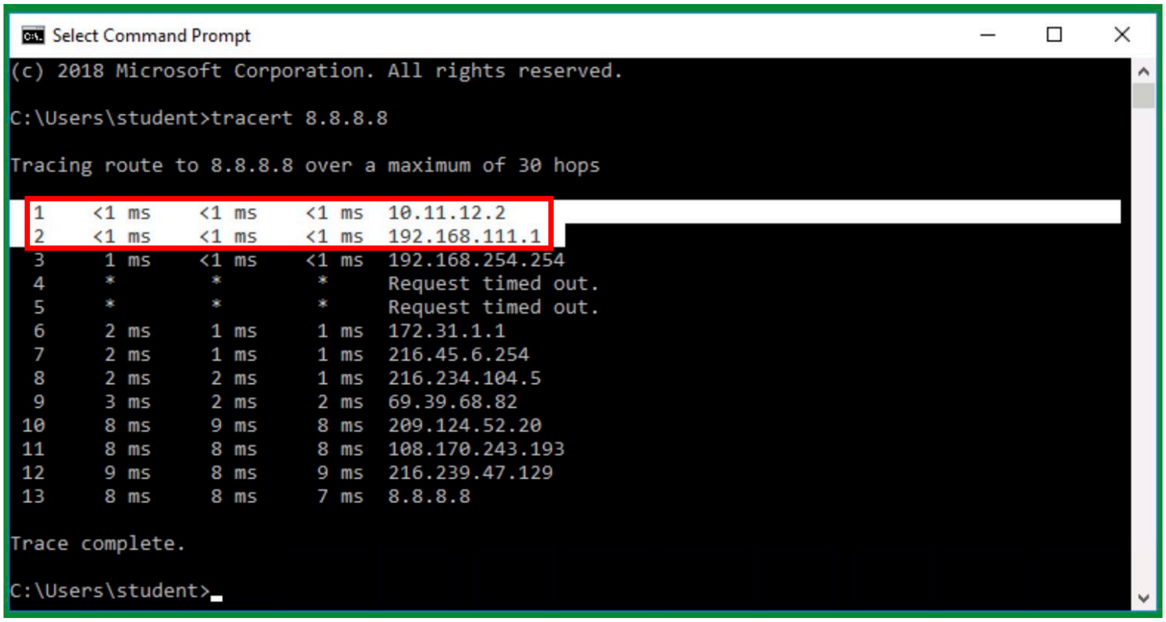

```
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.12.2
  2    <1 ms    <1 ms    <1 ms  192.168.111.1
  3     1 ms    <1 ms    <1 ms  192.168.254.254
  4     *        *        *     Request timed out.
  5     *        *        *     Request timed out.
  6     2 ms     1 ms     1 ms  172.31.1.1
  7     2 ms     1 ms     1 ms  216.45.6.254
  8     2 ms     2 ms     1 ms  216.234.104.5
  9     3 ms     2 ms     2 ms  69.39.68.82
 10     8 ms     9 ms     8 ms  209.124.52.20
 11     8 ms     8 ms     8 ms  108.170.243.193
 12     9 ms     8 ms     9 ms  216.239.47.129
 13     8 ms     8 ms     7 ms  8.8.8.8

Trace complete.

C:\Users\student>
```

Figure 10-7: Traceroute

What are the first and second hops?

____

Are they the same as in step 14?

____

IMPORTANT: Traffic from users in VLAN X11 is using Core-1 as the gateway, who in turn uses ISP-1 as the next hop. Users in VLAN X12 use Core-2 as the gateway and ISP-2 as the next hop (see figure 10-8). This behavior provides a load balancing effect across both ISPs. It leverages the customer's two services.

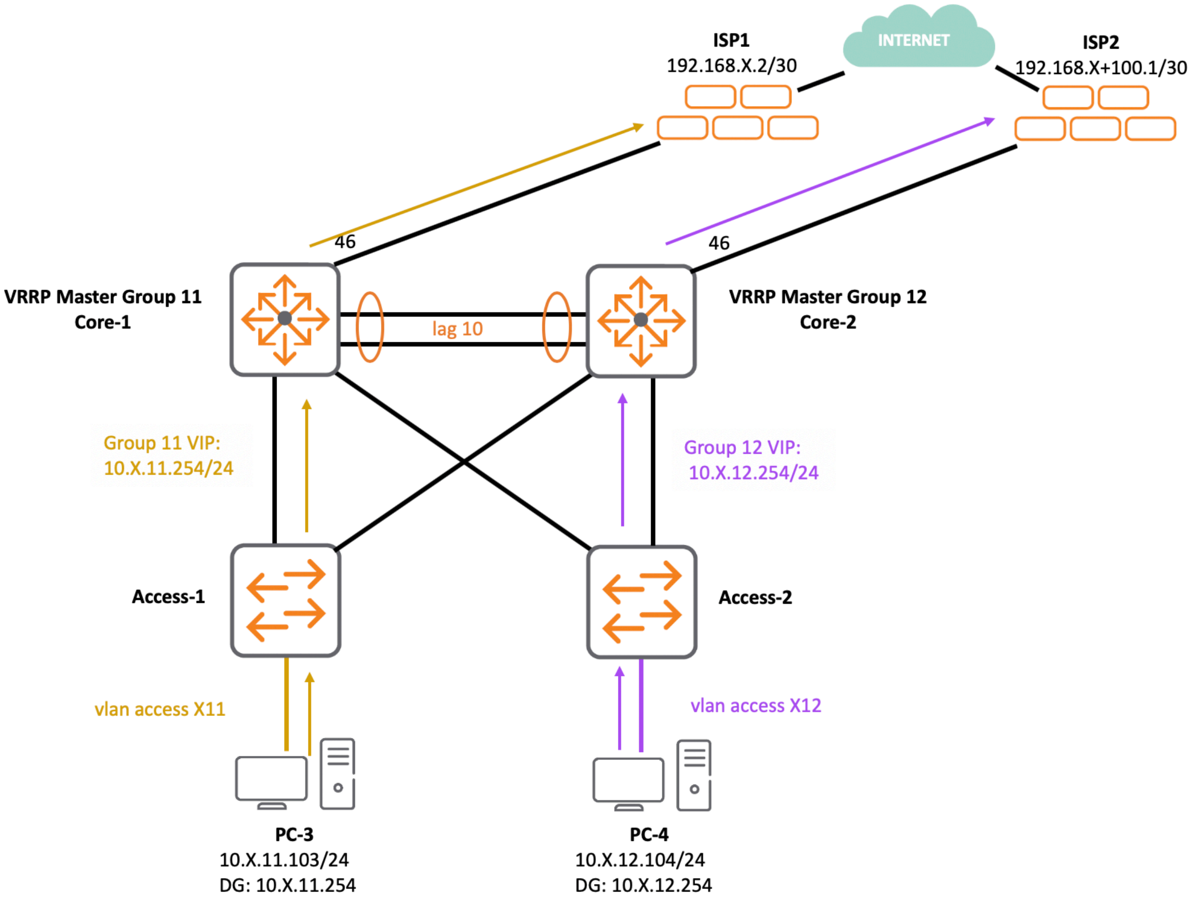

Figure 10-8: Traffic path

Task 3: Redundancy with Floating Routes

Objectives

Your current deployment has proven more efficient, however, it still has a weak point - it contains single points of failure. If the link to ISP1 fails, then users in VLAN X11 lose internet access. A similar result would occur to VLAN X12 clients if ISP2 fails. The solution to this is the creation of static floating routes.

In this task, you will create a second prefix on each Core pointing to the other Core. However, these prefixes will have a lower preference because of an increased administrative distance. When the main internet link on either Core is active, then the floating routes are not present in the routing table and not used. However, if the connection to either carrier goes down, the main route vanishes and the floating route is inserted and makes the switch route data traffic through its neighbor.

Additionally, there will be a new IP segment used as a layer 3 transport between the Cores. You already calculated this segment in Lab 9.1 Task 4b (Subnetting and VLSM).

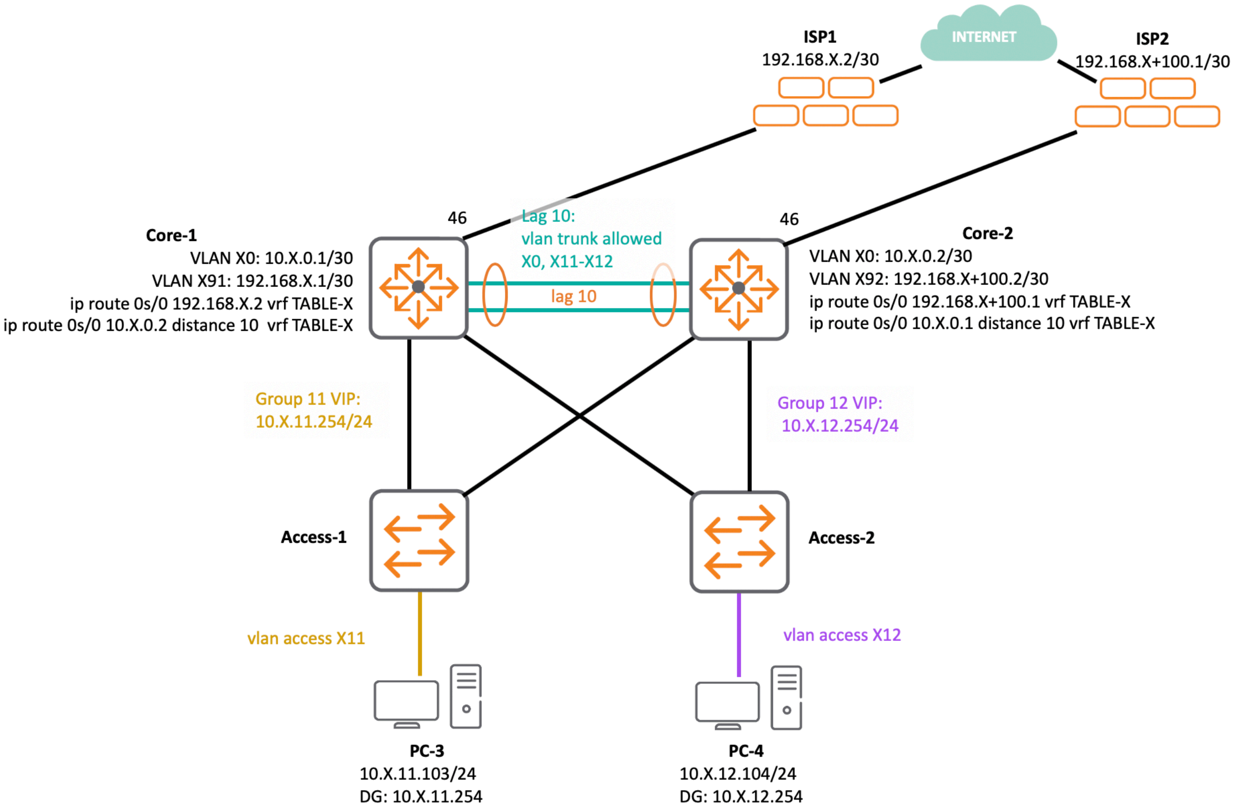

Figure 10-9: Task 3 topology

Steps

Core-1 (via PC-1)

1. Open a SSH session to Core-1.
2. Create **VLAN X0** and name it **CORE-1&2_TABLE-X**.

> **NOTE:** Replace the highlighted “X” for your student table number.

```
Core-1# configure terminal
Core-1(config)# vlan X0
Core-1(config-vlan-110)# name CORE-1&2_TABLE-X
Core-1(config-vlan-110)# exit
```

3. Allow VLAN X0 to LAG 10.

```
Core-1(config)# interface lag 10
Core-1(config-lag-if)# vlan trunk allow X0
Core-1(config-lag-if)# exit
```

4. Create interface **VLAN X0** and map it to vrf **TABLE-X**, then assign it the **10.X.0.1/30** IP address.

```
Core-1(config)# interface vlan X0
Core-1(config-if-vlan)# vrf attach TABLE-X
Core-1(config-if-vlan)# ip address 10.X.0.1/30
Core-1(config-if-vlan)# exit
```

5. Create a static default route in vrf **TABLE-X** pointing to **10.X.0.2** and assign it a **distance 10** (future Core-2 address in VLAN X0).

```
Core-1(config)# ip route 0.0.0.0/0 10.X.0.2 distance 10 vrf TABLE-X
Core-1(config)#
```

6. Show the static routes of vrf **TABLE-X**.

```
Core-1(config)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  192.168.11.2,  [1/0],  static

Core-1(config)#
```

How many entries do you have?

____

What is the next hop?

____

Core-2 (via PC-1)

7. Open a SSH session to Core-2.
8. Repeat steps 2 to 6 assigning **10.X.0.2/30** to Core-2 and use **10.X.0.1** as the route's next hop.

```
Core-2# configure terminal
Core-2(config)# VLAN X0
Core-2(config-vlan-110)# name CORE-1&2_TABLE-X
Core-2(config-vlan-110)# exit
```

```
Core-2(config)# interface lag 10
Core-2(config-lag-if)# vlan trunk allow X0
Core-2(config-lag-if)# exit
```

```
Core-2(config)# interface vlan X0
Core-2(config-if-vlan)# vrf attach TABLE-X
Core-2(config-if-vlan)# ip address 10.X.0.2/30
Core-2(config-if-vlan)# exit
```

```
Core-2(config)# ip route 0.0.0.0/0 10.X.0.1 distance 10 vrf TABLE-X
Core-2(config)#
```

```
Core-2(config)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  192.168.111.1,  [1/0],  static

Core-2(config)# exit
```

NOTICE: Creating identical routes on two layer 3 devices pointing to each other may lead to layer 3 loops. In our scenario, that would occur if both ISP links go down. In this unlikely case, if Core-1 receives traffic to the internet it would use Core-2 as next hop, who, in absence of its main internet link, would then send traffic back to Core-1, who would repeat the same process over and over and over.

Although there is a built-in layer 3 loop attenuation mechanism in the IP header, Time to Live (TTL), monitoring the validity of the floating route through Service Level Agreements* (SLAs) based tracking is always recommended in order to prevent this issue from happening, otherwise loop packets would consume data plane resources before they die.

* SLAs are out of the scope of this training.

TIP: As alternative to floating routes you can combine static routes along with either BGP conditional advertisement or IGPs default route injection. This approach prevents layer 3 loops entirely. You will examine IGP default route injection approach in the next lab.

In this part of the process, routes and traffic paths remain as they were in the end of task 2. You will now simulate a failure and confirm the resulting path.

PC-3 and PC-4

9. Access both PCs.

10. Run a continuous ping towards **8.8.8.8**. Pings should be successful.

Core-1 (via PC-1)

11. Move back to Core-1.
12. Disable interface VLAN X91.

```
Core-1(config)# interface vlan X91
Core-1(config-if-vlan)# shutdown
Core-1(config-if-vlan)#
```

13. Display the vrf **TABLE-X** routing table.

```
Core-1(config-if-vlan)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  10.11.0.2,  [10/0],  static

Core-1(config-if-vlan)#
```

What is the next hop of the 0's prefix now?

PC-3

14. Move to PC-3**.**

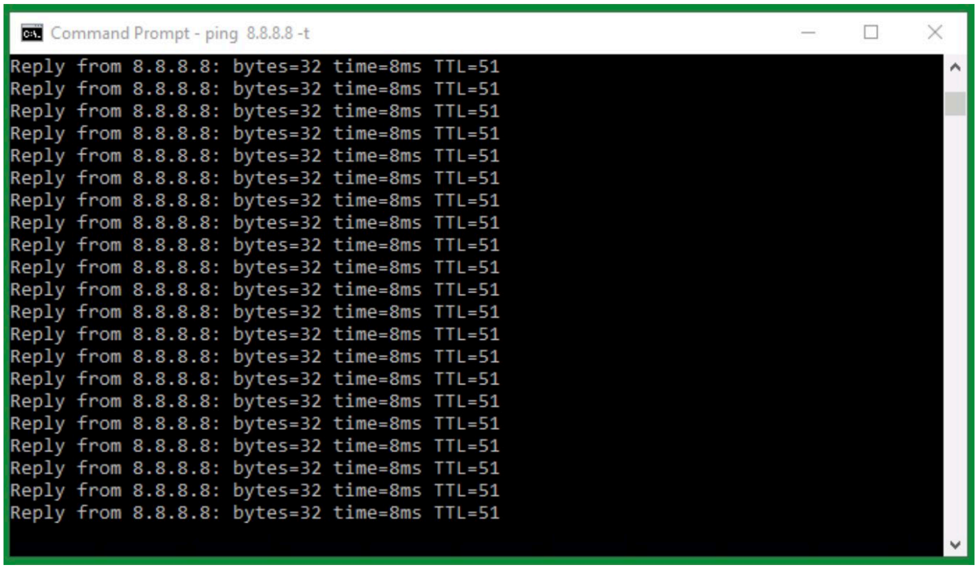

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
```

Figure 10-10: Ping -t

What is the ping status?

____

15. Run a traceroute towards **8.8.8.8.**

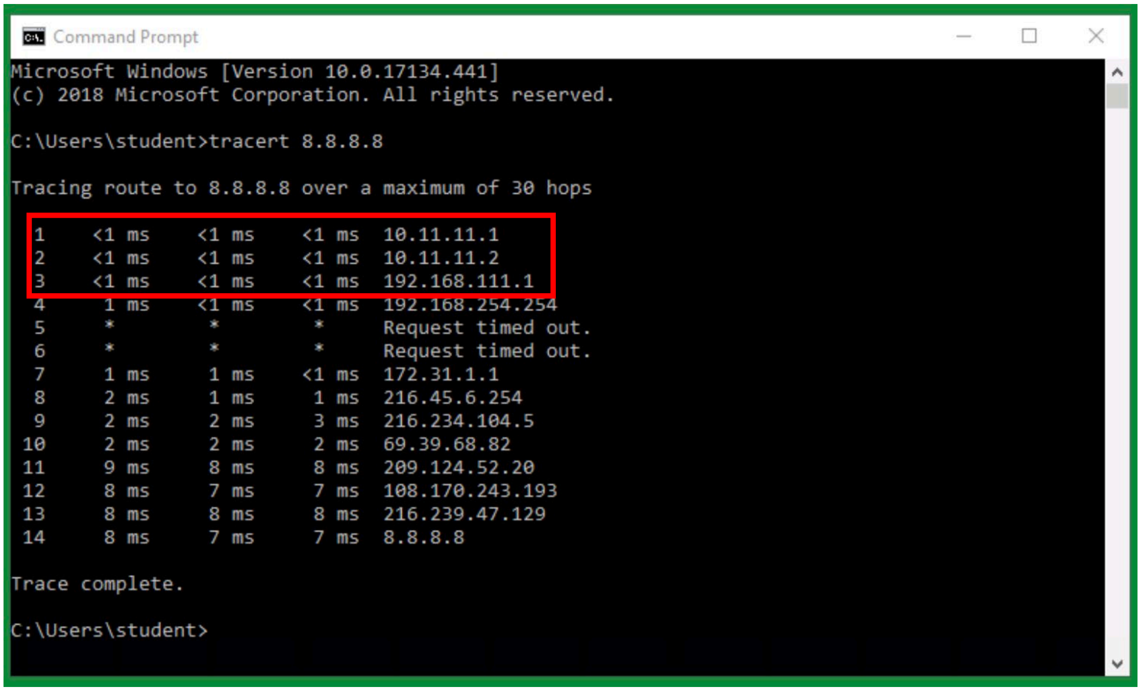

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.11.2
  3    <1 ms    <1 ms    <1 ms  192.168.111.1
  4     1 ms    <1 ms    <1 ms  192.168.254.254
  5     *        *        *     Request timed out.
  6     *        *        *     Request timed out.
  7     1 ms     1 ms    <1 ms  172.31.1.1
  8     2 ms     1 ms     1 ms  216.45.6.254
  9     2 ms     2 ms     3 ms  216.234.104.5
 10     2 ms     2 ms     2 ms  69.39.68.82
 11     9 ms     8 ms     8 ms  209.124.52.20
 12     8 ms     7 ms     7 ms  108.170.243.193
 13     8 ms     8 ms     8 ms  216.239.47.129
 14     8 ms     7 ms     7 ms  8.8.8.8

Trace complete.

C:\Users\student>
```

Figure 10-11: Traceroute

What are the first 3 hops?

1st hop: ____

2nd hop: ____

3rd hop:____

PC-4

16. Move to PC-4 then repeat step 15 (ping 8.8.8.8).

What is the ping status?

____

What are the first 2 hops?

1st hop: ____

2nd hop: ____

You have successfully deployed internet access redundancy and made BigStartup network resilient to failures, as shown in figure 10-12.

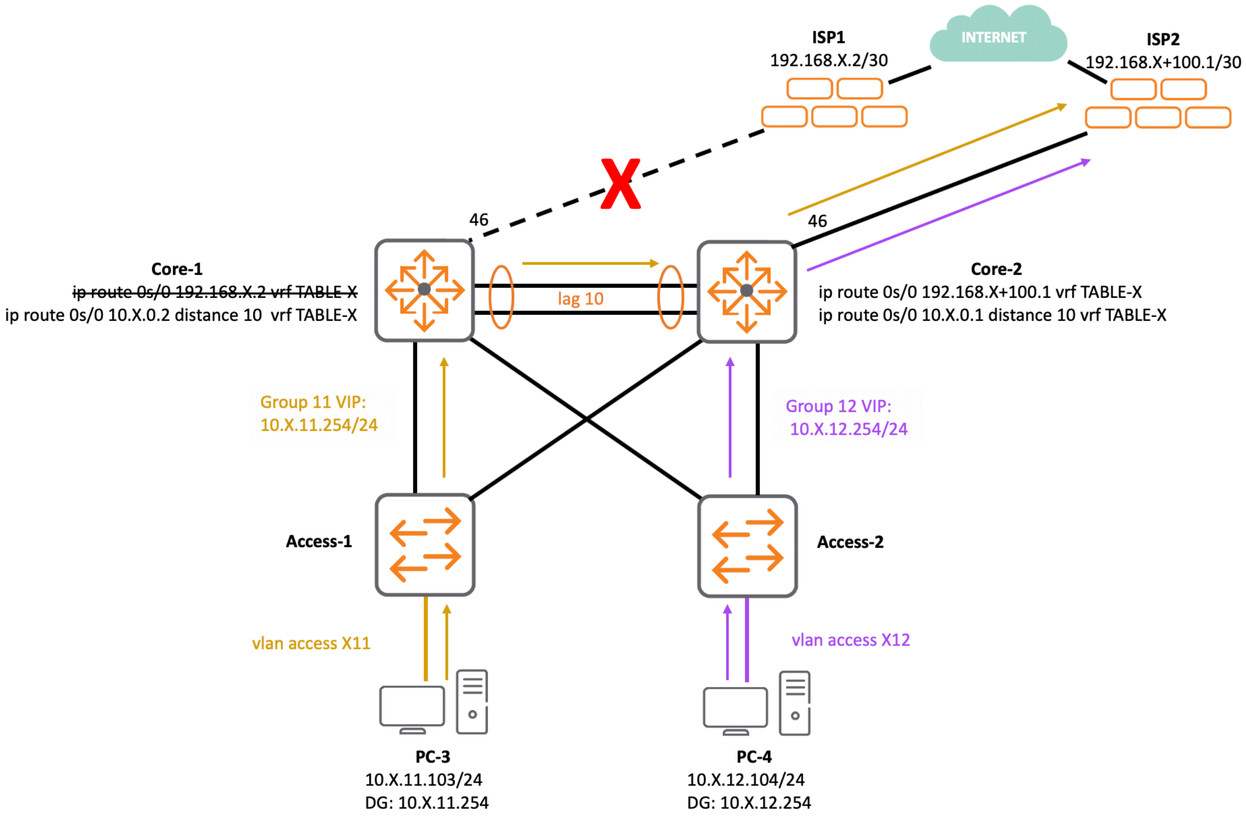

Figure 10-12: Floating routes failover

NOTE: Pay attention to Step 15: PC-3's second hop. Since Core-1 is delivering the packet to Core-2 on VLAN X0, then you would normally expect that hop 2 should be 10.X.0.2, however it is not. The logic behind this behavior is that when Core-2 receives from Core-1 the ICMP echo with TTL=1 then it subtracts 1, TTL becomes 0 and the packet dies as normal. Here Core-2 needs to respond back to the source (PC-3) with an "ICMP Time Exceeded" message (which is what you see on the tracert command's output), however according to Core-2's routing table, PC-3's IP address (10.X.11.103) isn't reachable via Core-1 on VLAN X0, but via VLAN X11 as a connected network (see output below), therefore it delivers the packet using layer 2. It uses the address it has from VLAN X11. This is called asymmetric routing (figure 10-13).

```
Core-2# show ip route 10.11.11.103 vrf TABLE-11
Displaying ipv4 routes selected for forwarding
'[x/y]' denotes [distance/metric]
10.11.11.0/24, vrf TABLE-11
        via  vlan1111,  [0/0],  connected
Core-2#
```

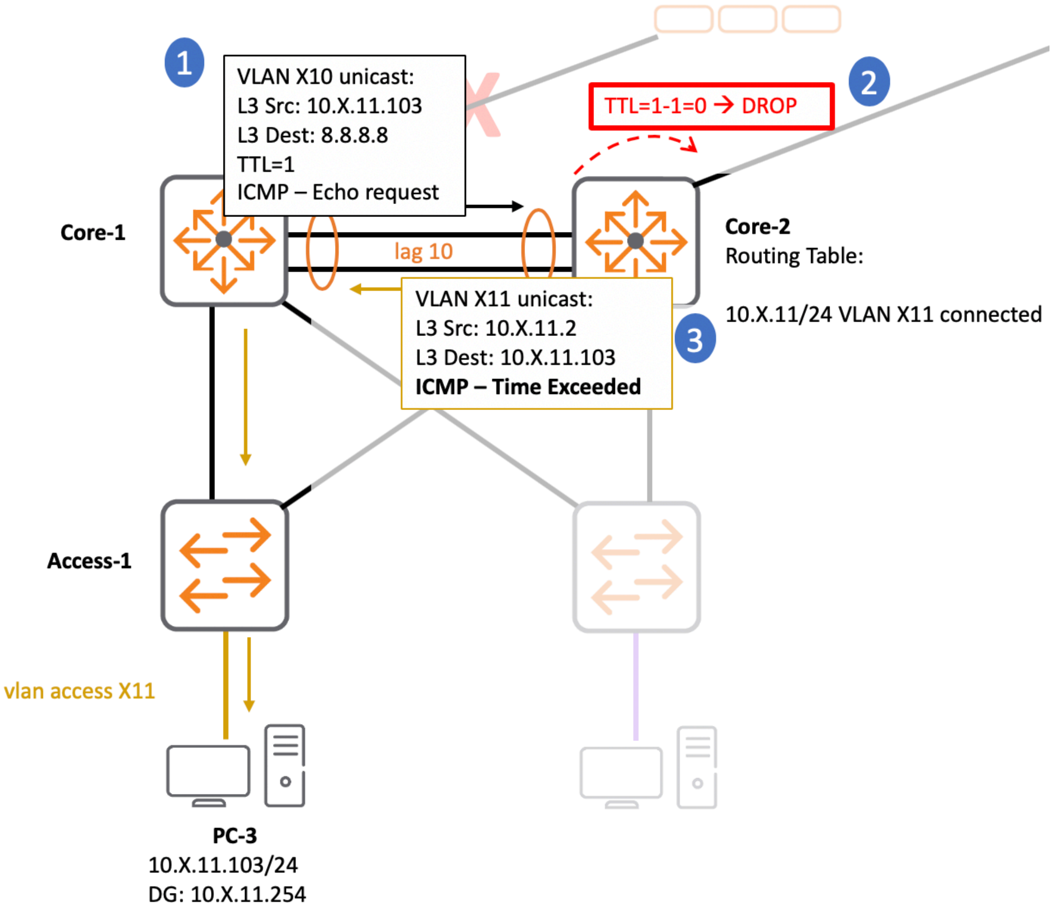

Figure 10-13: Asymmetric Routing

Core-1 (via PC-1)

17. Move back to Core-1.
18. Enable interface VLAN X91.

```
Core-1(config-if-vlan)# interface vlan X91
Core-1(config-if-vlan)# no shutdown
Core-1(config-if-vlan)# end
```

Task 4: Layer 3 Loop (optional).

Objectives

Creating identical routes on two layer 3 devices pointing to each other may lead to layer 3 loops. In our scenario, that would occur if both ISP links go down. In this unlikely case, if Core-1 receives traffic to the internet it would use Core-2 as next hop, who, in absence of its main internet link, would then send traffic back to Core-1, who would repeat the same process over and over and over again.

In this task you will simulate a failure on both internet connections in order to create a layer 3 loop and use the traceroute tool in order to diagnose it.

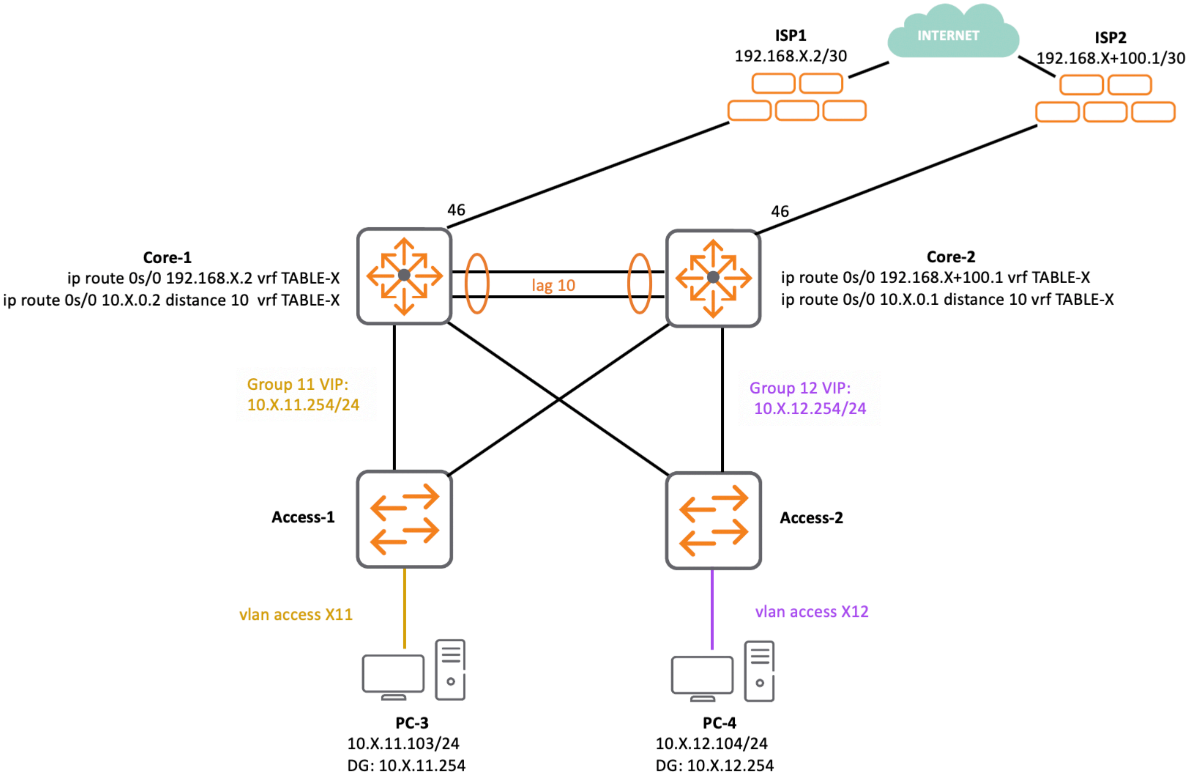

Figure 10-14: Task 4 Topology

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1.
2. Disable interface VLAN X91.

```
Core-1# configure terminal
Core-1(config)# interface vlan X91
Core-1(config-if-vlan)# shutdown
Core-1(config-if-vlan)#
```

3. Display Core-1's routing table.

```
Core-1(config-if-vlan)# show ip route static vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  10.11.0.2,  [10/0],  static

Core-1(config-if-vlan)#
```

Who is the next-hop for the 0s route?

ANSWER: Core-2.

Core-2 (via PC-1)

4. Open the SSH session to Core-2.
5. Disable interface VLAN X92.

```
Core-2# configure terminal
Core-2(config)# interface vlan X92
Core-2(config-if-vlan)# shutdown
Core-2(config-if-vlan)#
```

6. Display Core-2's routing table.

```
Core-2(config-if-vlan)# show ip route static vrf TABLE-X
Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  10.11.0.1,  [10/0],  static

Core-2(config-if-vlan)#
```

Who is the next-hop for the 0s route?

____

ANSWER: Core-1.

PC-3

7. Open a console session to PC-3.
8. Run a ping to **8.8.8.8**.

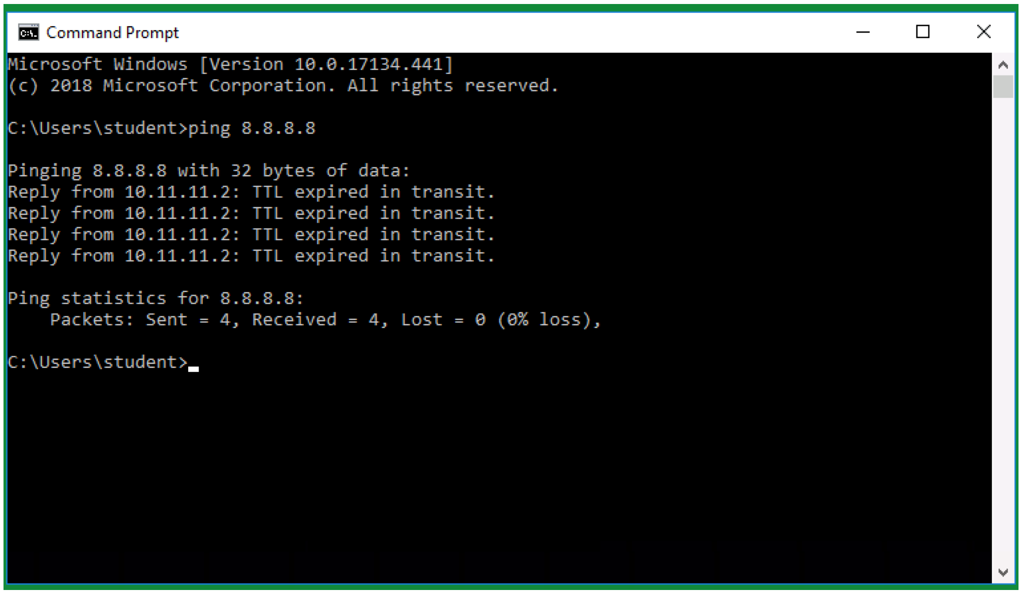

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 8.8.8.8

Pinging 8.8.8.8 with 32 bytes of data:
Reply from 10.11.11.2: TTL expired in transit.
Reply from 10.11.11.2: TTL expired in transit.
Reply from 10.11.11.2: TTL expired in transit.
Reply from 10.11.11.2: TTL expired in transit.

Ping statistics for 8.8.8.8:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\student>
```

Figure 10-15: Ping unsuccessful

Was the ping successful?

____

9. Run a traceroute to **8.8.8.8**.

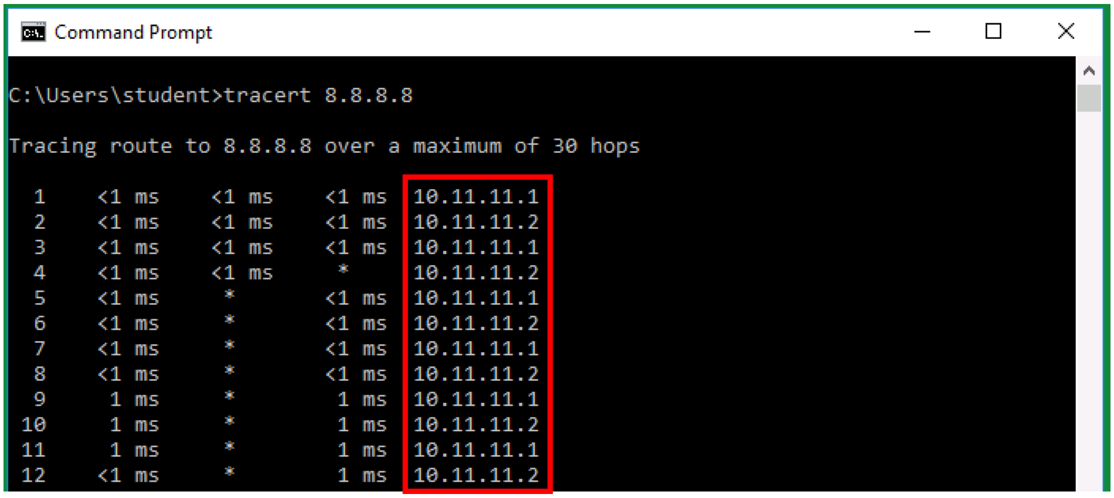

```
C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.11.2
  3    <1 ms    <1 ms    <1 ms  10.11.11.1
  4    <1 ms    <1 ms     *     10.11.11.2
  5    <1 ms     *       <1 ms  10.11.11.1
  6    <1 ms     *       <1 ms  10.11.11.2
  7    <1 ms     *       <1 ms  10.11.11.1
  8    <1 ms     *       <1 ms  10.11.11.2
  9     1 ms     *        1 ms  10.11.11.1
 10     1 ms     *        1 ms  10.11.11.2
 11     1 ms     *        1 ms  10.11.11.1
 12    <1 ms     *        1 ms  10.11.11.2
```

Figure 10-16: Traceroute Layer 3 loop

What are the next-hops shown in the output?

____

____

ANSWER: The output is showing Core-1 and Core-2 as the next-hops. The IP addresses in the output are the ones configured on VLAN X11. The reason for this is that at the time the packets die, the layer 3 switches return a time exceeded ICMP echo message back to PC-3. PC-3 is aware of both switches through interface VLAN X11.

Why does the output show traffic bouncing between the two Layer 3 devices?

____

ANSWER: Tracert works by increasing the TTL in a series of pings. When TTL = 1 the first hop is Core-1 (PC-3's gateway). When TTL = 2 the packet is sent to Core-1 who in turn gives it to its next-hop (Core-2). Next when TTL = 3, PC-3 gives the packet to Core-1 who in turn gives it to Core-2 as expected, however since Core-2 uses Core-1 as a its next-hop (since it lost its Internet connection as well) the packet comes back to Core-2. As seen in the output, traffic keeps bouncing between the two switches. This is a clear symptom of a Layer 3 loop, see figure 10-17.

Microsoft Windows will keep sending packets increasing the TTL one at the time, until it tries with TTL = 30, which is the last attempt it does. Other Operating Systems may use different thresholds.

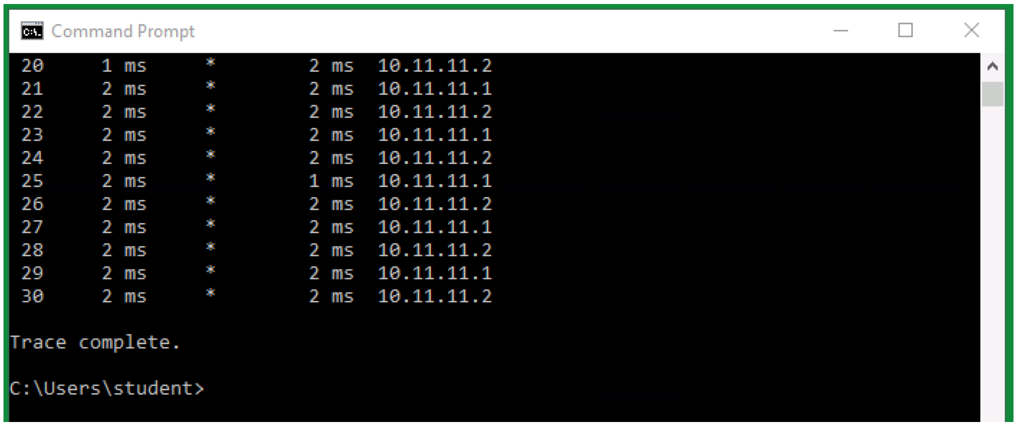

```
 20     1 ms     *        2 ms  10.11.11.2
 21     2 ms     *        2 ms  10.11.11.1
 22     2 ms     *        2 ms  10.11.11.2
 23     2 ms     *        2 ms  10.11.11.1
 24     2 ms     *        2 ms  10.11.11.2
 25     2 ms     *        1 ms  10.11.11.1
 26     2 ms     *        2 ms  10.11.11.2
 27     2 ms     *        2 ms  10.11.11.1
 28     2 ms     *        2 ms  10.11.11.2
 29     2 ms     *        2 ms  10.11.11.1
 30     2 ms     *        2 ms  10.11.11.2

Trace complete.

C:\Users\student>
```

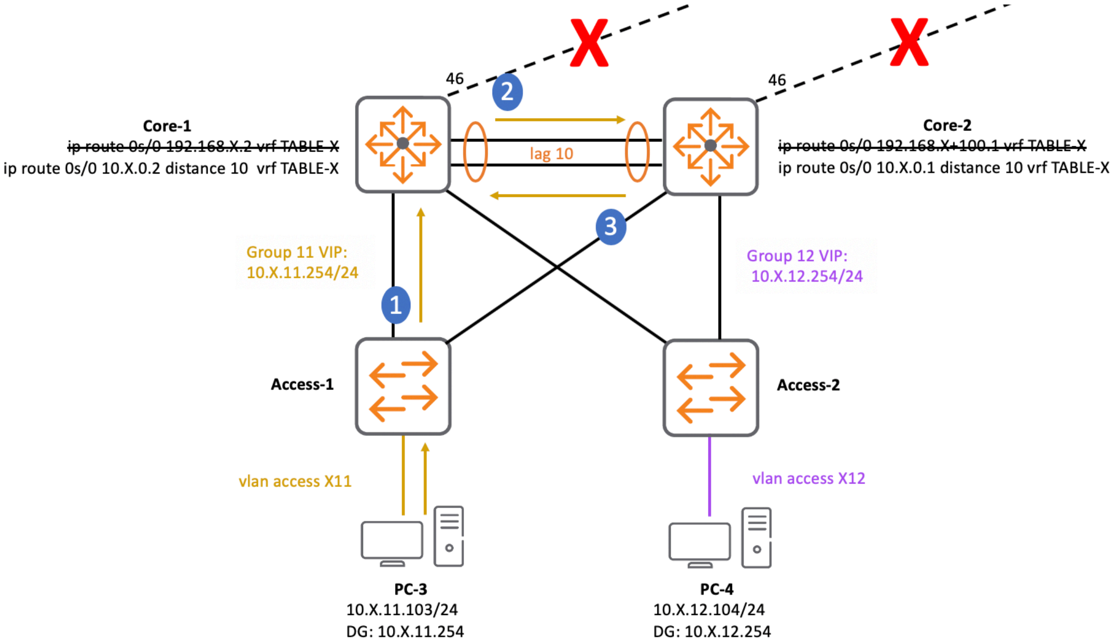

Figure 10-17: Layer 3 loop

IMPORTANT: Although there is a built-in layer 3 loop attenuation mechanism in the IP header called Time to Live (TTL); monitoring the validity of the floating route through Service Level Agreements* (SLAs) based tracking is always recommended in order to prevent this issue. Loop packets can consume data plane resources before they die.

* SLAs are out of the scope of this training.

TIP: As alternative to floating routes you can combine static routes along with either BGP conditional advertisement or IGPs default route injection. This approach prevents layer 3 loops entirely. You will examine the IGP default route injection approach in a following lab.

You will now fix the issue and see the effects while you do it.

10. Repeat the traceroute to **8.8.8.8.**

Core-1 (via PC-1)

11. Move to Core-1.
12. Enable interface VLAN X91 and

```
Core-1(config-if-vlan)# interface vlan X91
Core-1(config-if-vlan)# no shutdown
Core-1(config-if-vlan)# end
```

PC-3

13. Move back to PC-3. You will see traffic forwarded properly after bouncing between Core-1 and Core-2.

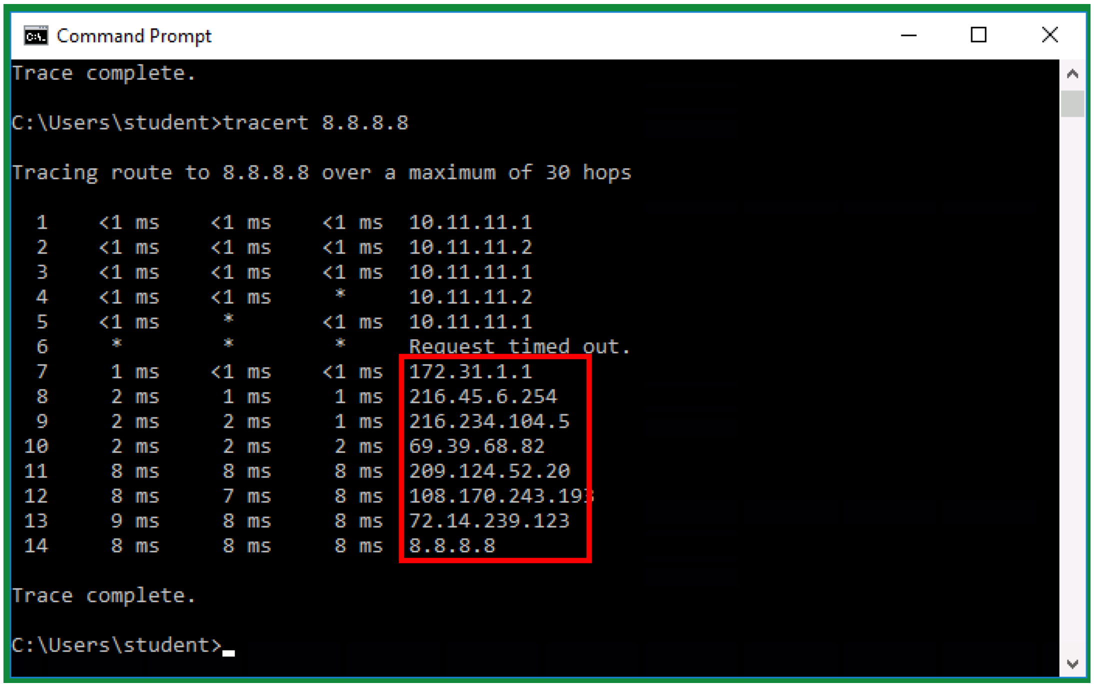

```
Trace complete.

C:\Users\student>tracert 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.11.2
  3    <1 ms    <1 ms    <1 ms  10.11.11.1
  4    <1 ms    <1 ms     *     10.11.11.2
  5    <1 ms     *       <1 ms  10.11.11.1
  6     *        *        *     Request timed out.
  7     1 ms    <1 ms    <1 ms  172.31.1.1
  8     2 ms     1 ms     1 ms  216.45.6.254
  9     2 ms     2 ms     1 ms  216.234.104.5
 10     2 ms     2 ms     2 ms  69.39.68.82
 11     8 ms     8 ms     8 ms  209.124.52.20
 12     8 ms     7 ms     8 ms  108.170.243.193
 13     9 ms     8 ms     8 ms  72.14.239.123
 14     8 ms     8 ms     8 ms  8.8.8.8

Trace complete.

C:\Users\student>_
```

Figure 10-18: Layer 3 loop recovery.

Core-2 (via PC-1)

1. Move to Core-2.
2. Enable VLAN X92 and

```
Core-2(config-if-vlan)# interface vlan X92
Core-2(config-if-vlan)# no shutdown
Core-2(config-if-vlan)# end
```

Task 5: Save Your Configurations

Objectives

You will now proceed to save your configuration.

Steps

Core-1 and Core-2 (via PC-1)

1. Save the current Cores’ configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

You have completed Lab 10!

AOS-CX Switching Fundamentals

Lab 11.1: Open Shortest Path First

Overview

This morning, while drinking your coffee and browsing your email, you notice a message from BigStartup titled: “PO: Professional Services – Server Switch Integration”. A few hours later, you meet your customer and find out that the servers they ordered months ago have finally arrived, along with a Data Center grade 8325 AOS-CX switch intended for connecting them. Although another supplier, called NetAmateur, will be in charge of that switch’s implementation, they want you to take care of the Core part.

They also have plans for expanding and extending the network to remote locations in the following years, and they will want these locations to be able to access the servers. You have advised them this is a good time to design and deploy a dynamic routing protocol called OSPF.

Objectives

After completing this lab, you will be able to:

- Define an OSPF router ID
- Create VRF specific OSPF process
- Create an Area and assign it to interfaces
- Build neighbor relationships
- Validate OSPF learned prefixes
- Deploy DHCP Helper role

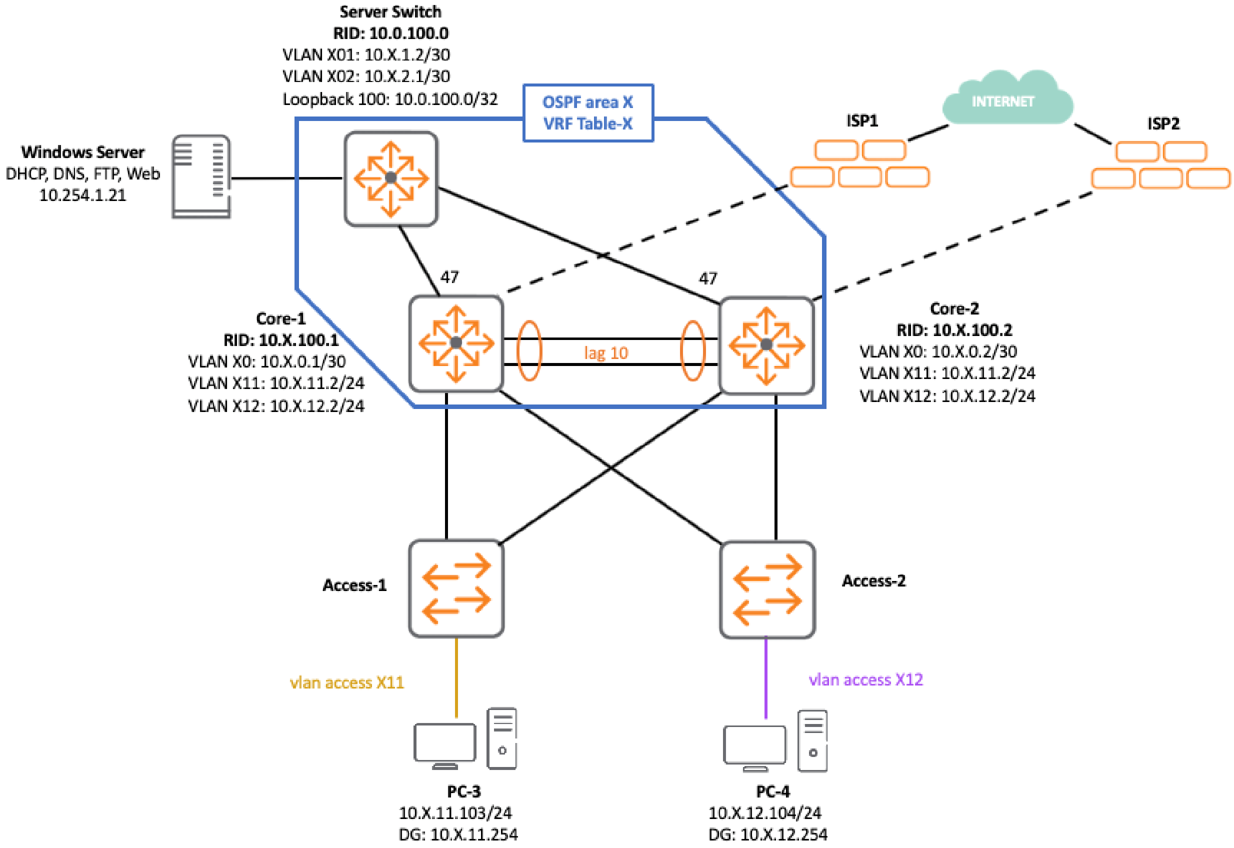

Figure 11.1-1: Lab Topology

Task 1: OSPF Single Area Between Cores

Objectives

You are about to run an OSPF single Area deployment on your core switches. This includes defining a unique Router ID, enabling the process and mapping it to a VRF, creating an OSPF area and assigning it to interfaces. You will begin with the link between Cores.

Once the tasks are completed you will proceed with neighbor discovery validation.

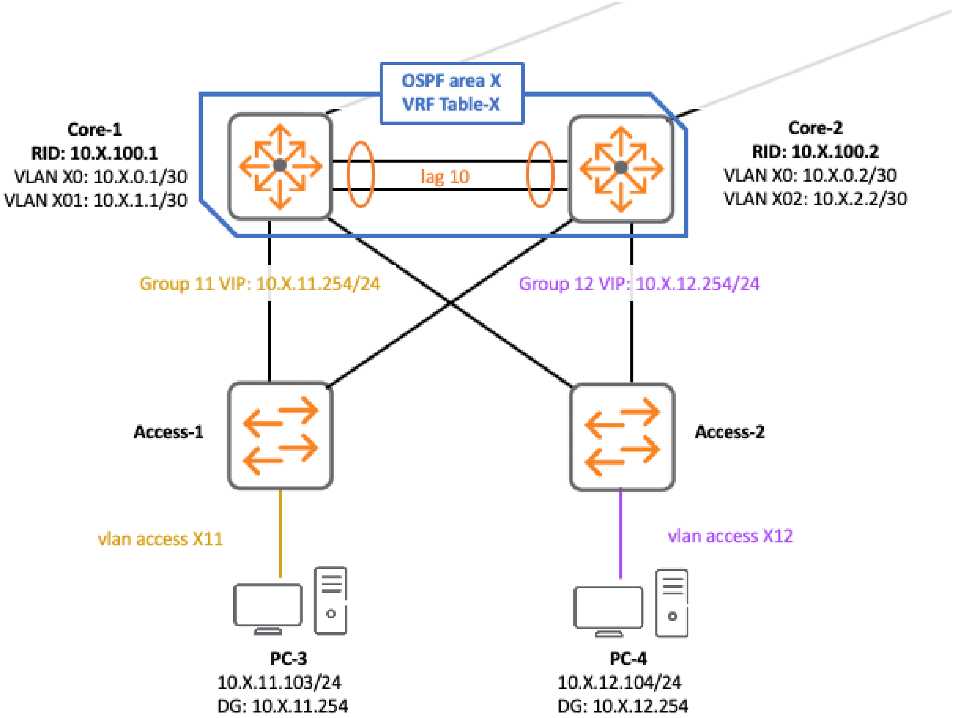

Figure 11.1-2: Task 1 topology

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1. Login using **cxfX/aruba123**

NOTE: Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

2. Define the height of the page to 40 lines.

```
Core-1# page 40
Core-1#
```

3. Create the OSPF process number X and map it to vrf **TABLE-X.**

NOTE: Replace the highlighted "X" for your student table number.

```
Core-1# configure terminal
Core-1(config)# router ospf X vrf TABLE-X
```

4. Assign **Router ID 10.X.100.1** and create **area X**.

```
Core-1(config-ospf-11)# router-id 10.X.100.1
Core-1(config-ospf-11)# area X
```

5. Enable the process.

```
Core-1(config-ospf-11)# enable
Core-1(config-ospf-11)# exit
```

NOTE: At this point OSPF is up and running in Core-1, however it is not sending Hello messages yet because you haven't enabled it on any interfaces. You will now enable it on the link to Core-2.

6. Assign OSPF process X area X on interface **VLAN X0**.

```
Core-1(config)# interface vlan X0
```

```
Core-1(config-if-vlan)# ip ospf X area X
Core-1(config-if-vlan)# end
```

7. Review the OSPF process state on vrf TABLE-X.

```
Core-1# show ip ospf vrf TABLE-X
Routing Process 11 with ID : 10.11.100.1 VRF TABLE-11
------------------------------------------------------

OSPFv2 Protocol is enabled
Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 NSSA
Number of active areas is 1, 1 normal, 0 stub, 0 NSSA
BFD is disabled
Reference Bandwidth: 100000 Mbps
Area (0.0.0.11) (Active)
  Interfaces in this Area: 1 Active Interfaces: 1
  Passive Interfaces: 0 Loopback Interfaces: 0
  SPF calculation has run 4 times
  Area ranges:
  Number of LSAs: 1, checksum sum 30066

Core-1#
```

What routing ID is this OSPF router using?

____

What is the state of the protocol?

____

How many areas are created and what is the Area ID?

____

How many LSAs have been created?

What LSA type do you think it is?

8. Display the status of OSPF interfaces on vrf **TABLE-X**.

```
Core-1# show ip ospf interface vrf TABLE-X
Interface vlan110 is up, line protocol is up
---------------------------------------------

IP address 10.11.0.1/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
    Transit delay 1 sec, Router priority 1
    Designated Router IP: 10.11.0.1
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled

Core-1#
```

What interface is running the protocol?

What area is the interface on?

What are the default Hello and Dead timers?

NOTE: Right now, Core-1 is sending hello messages out of Interface VLAN X0, however, there is no other OSPF router on that segment yet. You will proceed to deploy the counterpart on Core-2.

Core-2 (via PC-1)

9. Open the SSH session to Core-2.
10. Repeat steps 2 to 6 using **Router ID 10.X.100.2**.

```
Core-2# page 40
Core-2#
```

```
Core-2# configure terminal
Core-2(config)# router ospf X vrf TABLE-X
Core-2(config-ospf-11)# router-id 10.X.100.2
Core-2(config-ospf-11)# area X
Core-2(config-ospf-11)# enable
Core-2(config-ospf-11)# exit
```

```
Core-2(config)# interface vlan X0
Core-2(config-if-vlan)# ip ospf X area X
Core-2(config-if-vlan)# end
```

11. List all OSPF neighbors that Core-2 has discovered. Include the details.

```
Core-2# show ip ospf neighbors detail vrf TABLE-X
Neighbor 10.11.100.1, interface address 10.11.0.1
-------------------------------------------------

Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface vlan110
Neighbor priority is 1, State is FULL
DR is 10.11.0.1, BDR is 10.11.0.2
Options is 0x42
Dead timer due in 00:00:30
Retransmission queue length 0
Time since last state change 00h:02m:39s
Core-2#
```

What neighbor has Core-2 discovered? List the Router ID.

____

What is the adjacency state?

____

What is the Designated Router and what is the Backup?

____

Since both nodes have a default priority of 1, how was Designated Router (DR) elected?

____

12. Write down the roles for the links between the cores in figure 11.1-3.

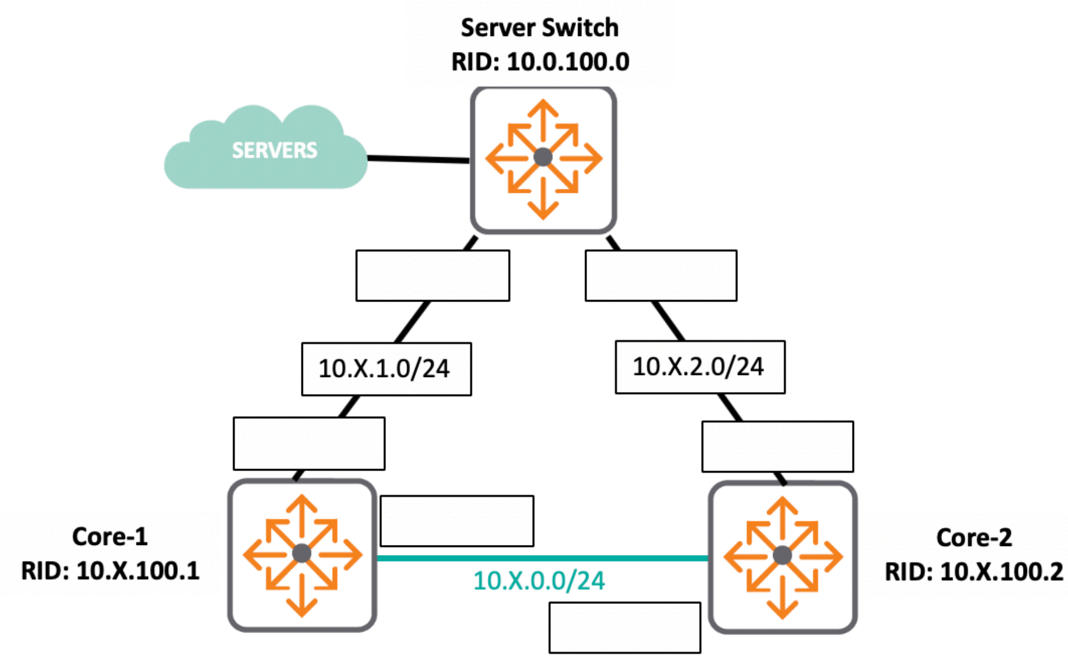

Figure 11.1-3: DRs and BDRs

TIP: You can find a larger copy of this diagram in Appendix 3.

13. Display the Router LSAs that Core-2 knows.

```
Core-2# show ip ospf X lsdb vrf TABLE-X router
OSPF Router with ID (10.11.100.2) (Process ID 11 VRF TABLE-11)
===============================================================

Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum      Link Count
------------------------------------------------------------------------------
10.11.100.1     10.11.100.1     1318      0x80000006 0x00009138      1
10.11.100.2     10.11.100.2     1324      0x80000006 0x00008f37      1

Core-2#
```

How many Router LSAs are shown?

What are the Link State IDs?

Who created those LSAs?

What information do they contain?

14. Display the Network LSAs that Core-2 knows.

```
Core-2# show ip ospf X lsdb vrf TABLE-X network
OSPF Router with ID (10.11.100.2) (Process ID 11 VRF TABLE-11)
==============================================================

Network Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.11.0.1       10.11.100.1     1479      0x80000001 0x00003b9b

Core-2#
```

How many Router LSAs are shown?

____

What is the Link State IDs?

____

Who created those LSAs?

____

What information do they contain?

____

Core-1 (via PC-1)

15. Move to Core-1.
16. Display the Router LSAs

```
Core-1# show ip ospf X lsdb vrf TABLE-X router
OSPF Router with ID (10.11.100.1) (Process ID 11 VRF TABLE-11)
```

```
===============================================================

Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.11.100.1     10.11.100.1     1318      0x80000006 0x00009934      1
10.11.100.2     10.11.100.2     1324      0x80000006 0x00009733      1

Core-1#
```

Are these LSAs similar to the ones that Core-2 has?

____

TIP: In order to confirm if they are the same version, you have to compare the LSID and Sequence number.

How many links do each of them announce?

____

NOTE: Right now, only one link is contained within the Router LSA (10.X.0.0/30).

Task 2: Add the Server Switch

Objectives

The next phase in this integration will be to build the interconnection with the Server Switch using the two links that have already been plugged in. You will use VLANs X01 and X02 for that as shown in figure 11.1-4.

Remember that you will only take care of Core switches configuration, while the Server Switch is being configured by another partner.

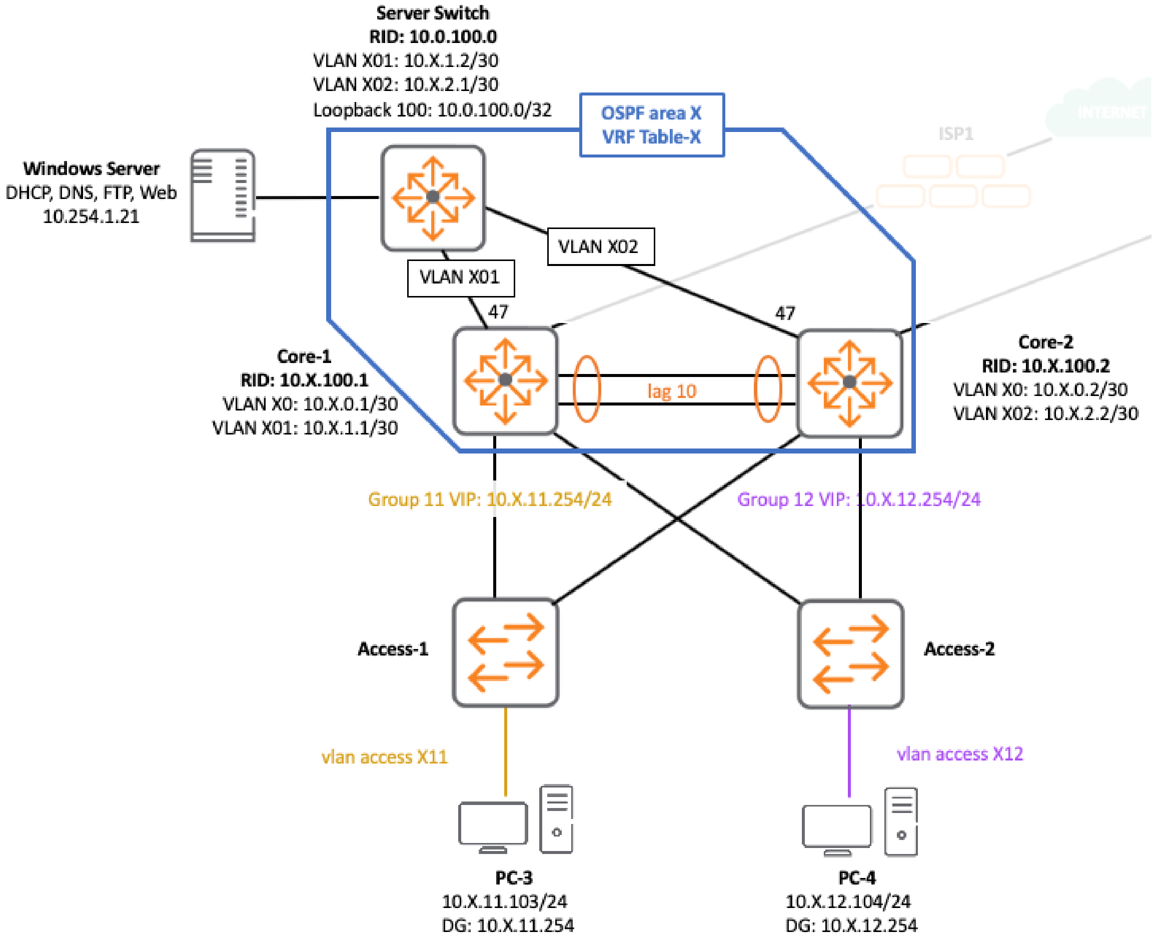

Figure 11.1-4: Task 2 topology

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1-
2. Create **VLAN X01** and put a name. This will be the transport for the uplink to the Server Switch.

```
Core-1# configure terminal
Core-1(config)# vlan X01
Core-1(config-vlan-1101)# name TO_SERVER-SWITCH_TABLE-X
Core-1(config-vlan-1101)# exit
```

3. Tag it on port **1/1/47**.

```
Core-1(config)# interface 1/1/47
Core-1(config-if)# vlan trunk allowed X01
Core-1(config-if)# exit
```

4. Create interface **VLAN X01** and map it to vrf **TABLE-11**, put the **10.X.1.1/30 IP address** and enable the OSPF process on it.

```
Core-1(config)# interface vlan X01
Core-1(config-if-vlan)# vrf attach TABLE-X
Core-1(config-if-vlan)# ip address 10.X.1.1/30
Core-1(config-if-vlan)# ip ospf X area X
Core-1(config-if-vlan)# end
```

5. Display the list of neighbors Core-1 has now.

```
Core-1# show ip ospf neighbors detail vrf TABLE-11
Neighbor 10.11.100.2, interface address 10.11.0.2
------------------------------------------------

Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface vlan110
Neighbor priority is 1, State is FULL
DR is 10.11.0.1, BDR is 10.11.0.2
Options is 0x42
Dead timer due in 00:00:36
```

```
Retransmission queue length 0
Time since last state change 00h:03m:11s
Neighbor 10.0.100.0, interface address 10.11.1.2
------------------------------------------------

Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface vlan1101
Neighbor priority is 1, State is FULL
DR is 10.11.1.2, BDR is 10.11.1.1
Options is 0x42
Dead timer due in 00:00:31
Retransmission queue length 0
Time since last state change 00h:00m:08s
Core-1#
```

Did you discover any new neighbors?

Who?

6. Write down the role for the links between Core-1 and the Server Switch in figure 11.1-3.

Core-2 (via PC-1)

7. Move to Core-2.
8. Repeat steps 2 to 4 using **VLAN X02** and **10.11.X.2/30** for the IP address.

```
Core-2(config)# vlan X02
Core-2(config-vlan-112)# name TO_SERVER-SWITCH_TABLE-X
Core-2(config-vlan-112)# exit
```

```
Core-2(config)# interface 1/1/47
Core-2(config-if)# vlan trunk allowed X02
Core-2(config-if)# exit
```

```
Core-2(config)# interface vlan X02
```

```
Core-2(config-if-vlan)# vrf attach TABLE-X
Core-2(config-if-vlan)# ip address 10.X.2.2/30
Core-2(config-if-vlan)# ip ospf X area X
Core-2(config-if-vlan)# end
```

9. Display the list of neighbors Core-2 has now.

```
Core-2# show ip ospf neighbors vrf TABLE-X
OSPF Process ID 11 VRF TABLE-11
================================

Total Number of Neighbors: 1

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
10.11.100.1      1         FULL/DR          10.11.0.1          vlan110

Core-2#
```

How many entries are listed?

____

Is there any device missing?

____

10. Confirm OSPF is properly enabled on interface **VLAN X02**.

```
Core-2# show ip ospf interface vlanX02 vrf TABLE-X
Interface vlan1102 is up, line protocol is up
---------------------------------------------

IP address 10.11.2.2/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
    Transit delay 1 sec, Router priority 1
    Designated Router IP: 10.11.2.2
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled
```

```
Core-2#
```

What is the status of the interface and protocol?

____

TIP: By looking at the configuration, it seems everything is in order. You will most likely have to look at packet statistics in order to see what packets are being exchanged between Core-2 and the Server switch.

11. Display the OSPF packet statistics for **interface VLAN X02.**

```
Core-2# show ip ospf statistics interface vlanX02 vrf TABLE-X
OSPF Process ID 11 VRF TABLE-11, interface vlan1102 statistics  (cleared 0h1m57s
ago)
=================================================================================
====

Tx Hello Packets      : 11          Rx Hello Packets      : 0
Tx Hello Bytes        : 704         Rx Hello Bytes        : 0
Tx DD Packets         : 0           Rx DD Packets         : 0
Tx DD Bytes           : 0           Rx DD Bytes           : 0
Tx LS Request Packets : 0           Rx LS Request Packets : 0
Tx LS Request Bytes   : 0           Rx LS Request Bytes   : 0
Tx LS Update Packets  : 0           Rx LS Update Packets  : 0
Tx LS Update Bytes    : 0           Rx LS Update Bytes    : 0
Tx LS Ack Packets     : 0           Rx LS Ack Packets     : 0
Tx LS Ack Bytes       : 0           Rx LS Ack Bytes       : 0

Total Number of State Changes : 6
Number of LSAs                : 0
LSA Checksum Sum              : 0
Total Transmit Failures       : 0
Total OSPF Packets Discarded  : 10

Reason                        Packets Dropped
-----------------------------------------------
Invalid type                  0
Invalid length                0
Invalid checksum              0
Invalid version               0
Bad or unknown source         0
Area mismatch                 0
Self-originated               0
Duplicate router ID           0
Interface standby             0
```

```
Total Hello packets dropped    5
  Network Mask mismatch        0
  Hello interval mismatch      5
  Dead interval mismatch       0
  Options mismatch             0
  MTU mismatch                 0
  Neighbor ignored             0
Authentication errors          0
  Type mismatch                0
  Authentication failures      0
Wrong protocol                 0
Resource failures              0
Bad LSA length                 0
Bad DD packets                 0
Others                         5

Total LSAs Ignored : 0
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0

Core-2#
```

Has Core-2 received any hello packet?

Has Core-2 dropped any hello packet?

Why?

ANSWER: Core-2 has dropped hello packets because of a hello interval mismatch. Although you know Core-2 is running the default value of 10 seconds, you are not certain of what interval value the Server Switch is using. You will have to run debugs in order to find out.

12. Clear the debug buffers.

```
Core-2# clear debug buffer
Core-2#
```

13. Display the ospfv2 debugs stored in buffers. This debug is on by default.

```
Core-2# show debug buff module ospfv2 | begin 2 10.0.100.0
2020-01-29:08:31:33.460518|hpe-routing|LOG_ERR|AMM|-|OSPFV2|OSPFV2|OSPF 268708866
Hello packet with mismatched hello interval received from router 10.0.100.0.
2020-01-29:08:31:33.460536|hpe-routing|LOG_ERR|AMM|-|OSPFV2|OSPFV2|My Hello
Interval = 10
2020-01-29:08:31:33.460550|hpe-routing|LOG_ERR|AMM|-|OSPFV2|OSPFV2|Neighboring
Hello Interval = 20
Core-2#
```

What information is the show debug displaying?

____

Is there any complaint about contents in hello messages?

____

NOTE: The output is clear, incoming hello packets interval is 2 times the usual one. Since this is a parameter that must match between two OSPF routers, the mismatch prevents the neighbor relationship from forming. When you share this information with the partner deploying the server switch (NetAmateur), you realize he is not an expert in the matter and does not understand what you are asking. However, he allows you to fix what you need to make this integration work.

Server Switch (via PC-1)

14. Using Putty, open an SSH session to the **Server-Switch**.

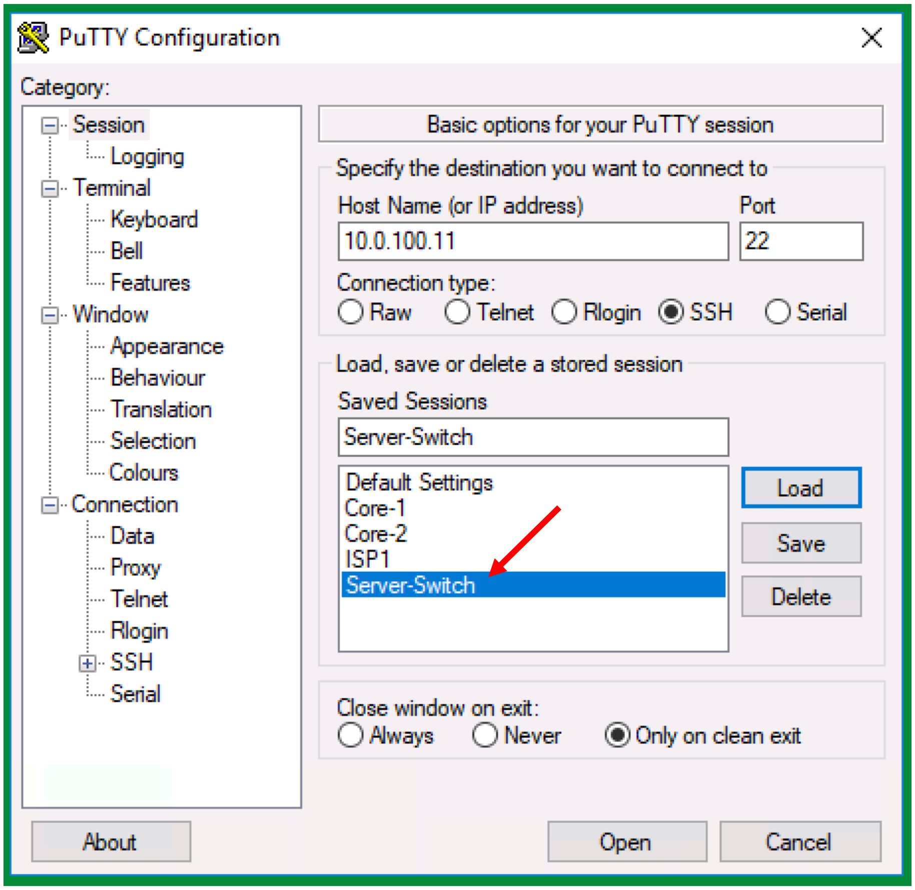

PuTTY Configuration

Category:

Basic options for your PuTTY session

Specify the destination you want to connect to

Host Name (or IP address): 10.0.100.11

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Load, save or delete a stored session

Saved Sessions: Server-Switch

Default Settings
Core-1
Core-2
ISP1
Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 11.1-5: SSH to Server Switch

15. Login using Username: **cxfX/aruba123.**

NOTE: Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

16. Validate current hello timer on **VLAN X02.**

```
ISP1-ServerSwitch# show ip ospf interface vlanX02 all-vrfs
Interface vlan102 is up, line protocol is up
---------------------------------------------

IP address 10.11.2.1/30, Process ID 1 VRF CXF_SERVER-SWITCH_TABLE-11, area
0.0.0.1
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
    Transit delay 1 sec, Router priority 1
    Designated Router IP: 10.11.2.1
    No backup designated router on this network
    Timer Intervals: Hello 20, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled

ISP1-ServerSwitch(config)#
```

What is the current hello timer value?

____

17. Decrease hello interval from 20 to 10 seconds on interface **VLAN X02**.

```
ISP1-ServerSwitch# configure terminal
ISP1-ServerSwitch(config)# interface vlan X02
ISP1-ServerSwitch(config-if-vlan)# ip ospf hello-interval 10
ISP1-ServerSwitch(config-if-vlan)# end
```

18. Close the putty session.

Core-2 (via PC-1)

19. Move back to Core-2.
20. Display the neighbors again. Server switch should be there.

```
Core-2# show ip ospf neighbors detail vrf TABLE-X
Neighbor 10.11.100.1, interface address 10.11.0.1
-------------------------------------------------

Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface vlan110
Neighbor priority is 1, State is FULL
DR is 10.11.0.1, BDR is 10.11.0.2
Options is 0x42
Dead timer due in 00:00:37
Retransmission queue length 0
Time since last state change 00h:15m:12s
Neighbor 10.0.100.0, interface address 10.11.2.1
------------------------------------------------

Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface vlan1102
```

```
Neighbor priority is 1, State is FULL
DR is 10.11.2.2, BDR is 10.11.2.1
Options is 0x42
Dead timer due in 00:00:35
Retransmission queue length 0
Time since last state change 00h:02m:04s
Core-2#
```

Did you discover any new neighbors?

____

Who?

____

21. Write down the role for the links between Core-1 and the Server Switch in figure 11.1-3.

NOTE: Right now, Core switches have each other and have the Server Switch as a neighbor. Therefore, they should be receiving Link State Updates that include the server's segment.

22. Display the routing table, including only the newly learned OSPF prefixes.

```
Core-2# show ip route ospf vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.0.100.0/32, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf
10.11.1.0/30, vrf TABLE-11
        via  10.11.0.1,  [110/200],  ospf
        via  10.11.2.1,  [110/200],  ospf
10.254.1.0/24, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf

Core-2#
```

What networks have Core-2 discovered?

____

____

____

What networks is the Server Switch a next hop for?

____

____

ANSWER: Server Switch (10.X.2.1) is the next hop for 10.0.100.0/32 and 10.254.1.0/24, its loopback and the servers' segment respectively.

What Is the Administrative Distance and metric for those segments?

____

Core-1 (via PC-1)

23. Move back to Core-1.
24. Display the routing table, including only the newly learned OSPF prefixes.

```
Core-1# show ip route ospf vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.0.100.0/32, vrf TABLE-11
        via  10.11.1.2,  [110/125],  ospf
10.11.2.0/30, vrf TABLE-11
        via  10.11.1.2,  [110/200],  ospf
        via  10.11.0.2,  [110/200],  ospf
10.254.1.0/24, vrf TABLE-11
        via  10.11.1.2,  [110/125],  ospf

Core-1#
```

What is the next-hop IP address for those networks learned from the server switch?

____

NOTE: Based on the outputs, both Cores are using their direct link to the server switch to reach segments that are beyond it (figure 11.1-5).

This also means that traffic from VLANs whose default gateway is Core-1 will be forwarded across VLAN X01, and traffic of VLANs whose default gateway is Core-2 will use the VLAN X02 uplink.

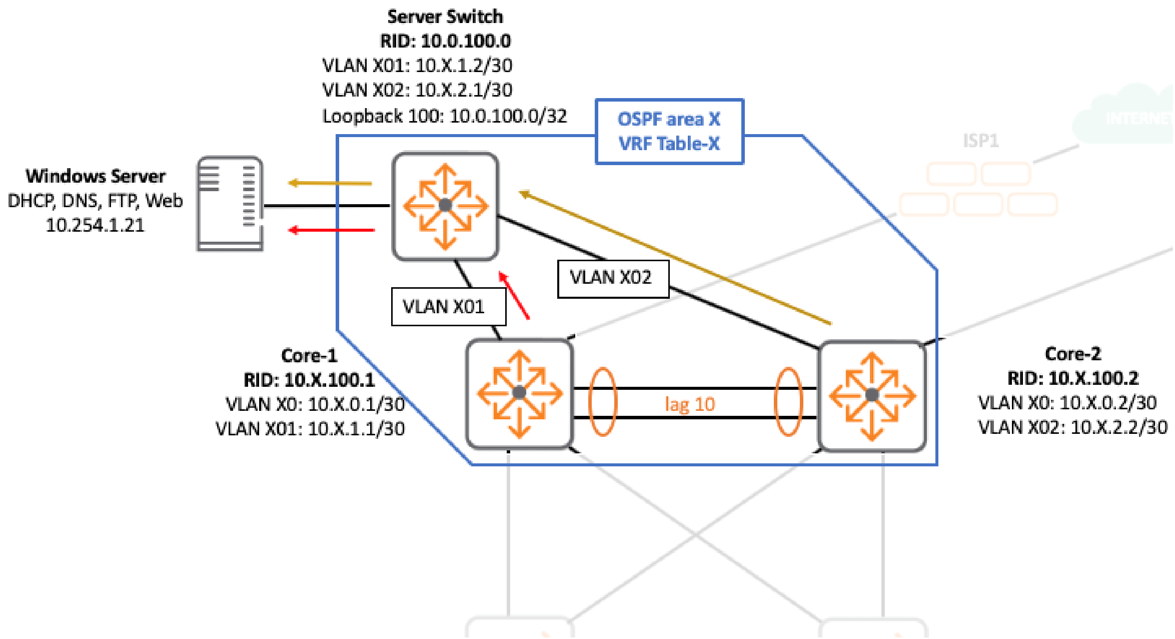

Figure 11.1-6: Upstream traffic

Now you will run some connectivity tests.

25. From Core-1, ping the Server Switch’s loopback IP address (**10.0.100.0**). The ping should be successful.

```
Core-1# ping 10.0.100.0 vrf TABLE-X
PING 10.0.100.0 (10.0.100.0) 100(128) bytes of data.
108 bytes from 10.0.100.0: icmp_seq=1 ttl=64 time=0.168 ms
108 bytes from 10.0.100.0: icmp_seq=2 ttl=64 time=0.164 ms
108 bytes from 10.0.100.0: icmp_seq=3 ttl=64 time=0.188 ms
108 bytes from 10.0.100.0: icmp_seq=4 ttl=64 time=0.166 ms
108 bytes from 10.0.100.0: icmp_seq=5 ttl=64 time=0.164 ms

--- 10.0.100.0 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4100ms
rtt min/avg/max/mdev = 0.164/0.170/0.188/0.009 ms
Core-1#
```

26. Ping the Windows Server IP address (**10.254.1.21**). The ping should be successful.

```
Core-1# ping 10.254.1.21 vrf TABLE-X
PING 10.254.1.21 (10.254.1.21) 100(128) bytes of data.
108 bytes from 10.254.1.21: icmp_seq=1 ttl=64 time=0.063 ms
108 bytes from 10.254.1.21: icmp_seq=2 ttl=64 time=0.066 ms
108 bytes from 10.254.1.21: icmp_seq=3 ttl=64 time=0.068 ms
108 bytes from 10.254.1.21: icmp_seq=4 ttl=64 time=0.068 ms
108 bytes from 10.254.1.21: icmp_seq=5 ttl=64 time=0.067 ms

--- 10.254.1.21 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4072ms
rtt min/avg/max/mdev = 0.063/0.066/0.068/0.007 ms
Core-1#
```

PC-3

27. Move to PC-3.

28. Ping the Windows Server (**10.254.1.21**).

Was ping successful?

____

29. Run a traceroute towards Windows Server.

Was it successful?

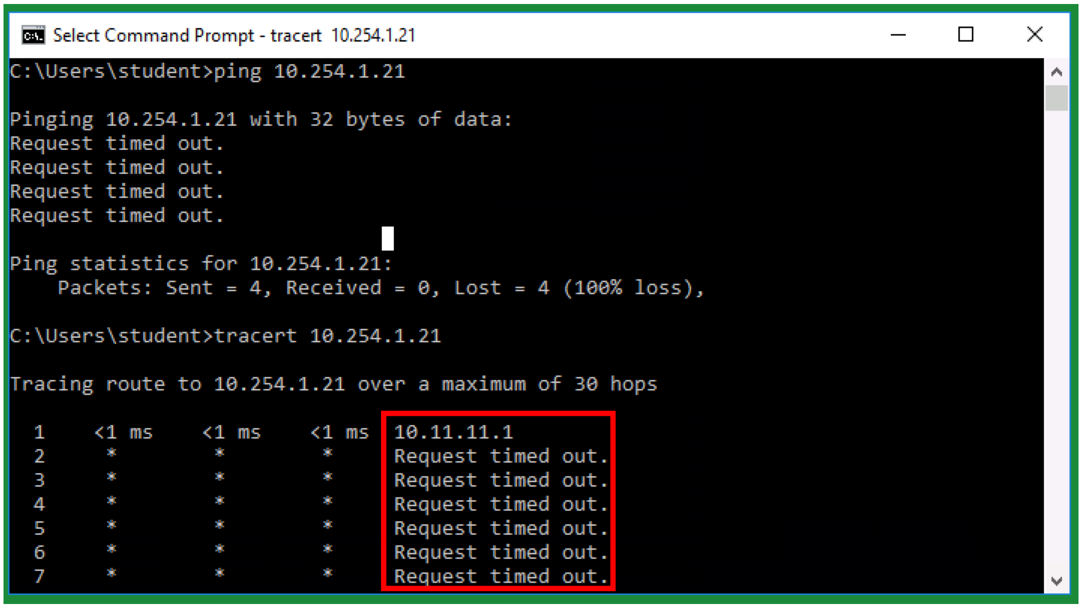

```
C:\Users\student>ping 10.254.1.21

Pinging 10.254.1.21 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 10.254.1.21:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2     *        *        *     Request timed out.
  3     *        *        *     Request timed out.
  4     *        *        *     Request timed out.
  5     *        *        *     Request timed out.
  6     *        *        *     Request timed out.
  7     *        *        *     Request timed out.
```

Figure 11.1-7: Failed traceroute

Why?

ANSWER: Traffic is failing for the same reason the first test to the Internet failed in lab 10. (Static Routes -Task 2 Step 7). Communications are bidirectional, it is not enough to know how to reach the remote destination but is also necessary that the other end knows how to send the replies back.

Task 3: Advertise LAN Segments

Objectives

In this activity you will advertise your LAN prefixes, so the Server Switch knows how to reach the client PCs.

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1.
2. Enable OSPF process X area X on interfaces **VLANs X11 - X12**.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11-X12
Core-1(config-if-vlan-<1111-1112>)# ip ospf X area X
Core-1(config-if-vlan-<1111-1112>)# end
```

3. Display the Router LSAs.

```
Core-1# show ip ospf X lsdb vrf TABLE-X router
OSPF Router with ID (10.11.100.1) (Process ID 11 VRF TABLE-11)
===============================================================

Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
------------------------------------------------------------------------------
10.0.100.0      10.0.100.0      73        0x80000013 0x0000d456     2
10.11.100.1     10.11.100.1     14        0x80000046 0x0000bc01     4
10.11.100.2     10.11.100.2     1737      0x8000004b 0x00006576     2

Core-1#
```

How many Router LSAs do you have now?

Who do they belong to?

How many links are counted for Core-1?

What links do they correspond to?

4. Confirm segments **10.X.11.0/24** and **10.X.12.0/24** are now part of the OSPF routing process.

```
Core-1# show ip ospf X routes vrf TABLE-X
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 11 VRF TABLE-11, Routing Table
-----------------------------------------------
```

```
Total Number of Routes : 8

10.0.100.0/32      (E1)
     via 10.11.1.2 interface vlan1101, cost 125 distance 110
10.11.0.0/30       (i) area: 0.0.0.11
     directly attached to interface vlan110, cost 100 distance 110
10.11.1.0/30       (i) area: 0.0.0.11
     directly attached to interface vlan1101, cost 100 distance 110
10.11.2.0/30       (i) area: 0.0.0.11
     via 10.11.0.2 interface vlan110, cost 200 distance 110
10.11.2.0/30       (i) area: 0.0.0.11
     via 10.11.1.2 interface vlan1101, cost 200 distance 110
10.11.11.0/24      (i) area: 0.0.0.11
     directly attached to interface vlan1111, cost 100 distance 110
10.11.12.0/24      (i) area: 0.0.0.11
     directly attached to interface vlan1112, cost 100 distance 110
10.254.1.0/24      (E1)
     via 10.11.1.2 interface vlan1101, cost 125 distance 110

Core-1#
```

Core-2 (via PC-1)

5. Move to Core-2.
6. Repeat step 2.

```
Core-2# configure terminal
Core-2(config)# interface vlan X11-X12
Core-2(config-if-vlan-<1111-1112>)# ip ospf X area X
Core-2(config-if-vlan-<1111-1112>)# end
Core-2#
```

Task 4: Testing Services

Objectives

In this activity you will start using one of the services that users in VLANs X11 and X12 have been waiting for: DHCP.

Since Layer 3 connectivity has been enabled all the way from the LAN segments up to the server's VLAN, then you can easily receive DHCP Discover messages at the Core switch and relay them up to the server. For redundancy, you will do it on both Cores.

Steps

Core-1 (via PC-1)

1. Open the SSH session to Core-1.
2. Move to interfaces VLAN X11 - X12, then enable the DHCP relay function.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11-X12
Core-1(config-if-vlan-<1111-1112>)# ip helper-address 10.254.1.21
Core-1(config-if-vlan-<1111-1112>)# end
```

Core-2 (via PC-1)

3. Open the SSH session to Core-2.
4. Repeat step 2.

```
Core-2# configure terminal
Core-2(config)# interface vlan X11-X12
Core-2(config-if-vlan-<1111-1112>)# ip helper-address 10.254.1.21
Core-2(config-if-vlan-<1111-1112>)# end
```

PC-3

5. Access PC-3.
6. Under the search field in the task bar, type "**control panel**". Windows will automatically display all items matching the string.
7. Click the top result (**Control Panel**). A new window will pop up.
8. In Control Panel, click "**View network status and tasks**" under Network and Internet.
9. Click "**Lab NIC**" under Access type Connections. A new window will pop up.
10. In Lab NIC status window, click "**Properties**" button.
11. In Lab NIC Properties section, select "**Internet Protocol Version 4 (TCP/IPv4)**, then click "**Properties**" button.
12. In Internet Protocol Version 4 (TCP/IPv4) Properties, choose "**Obtain an IP address automatically**" under General tab.
13. Then choose "**Obtain DNS server address automatically**".

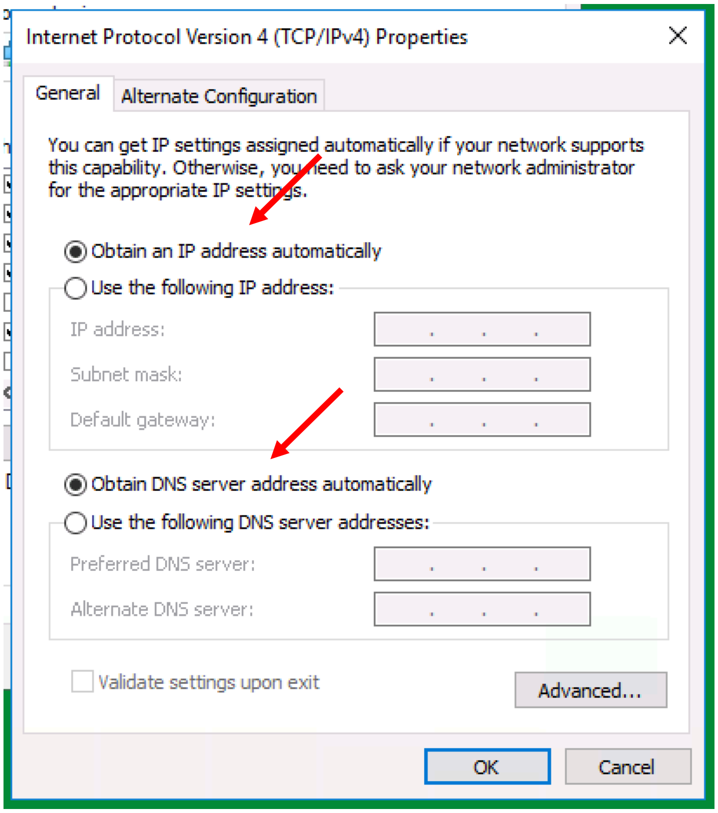

Internet Protocol Version 4 (TCP/IPv4) Properties

General | Alternate Configuration

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

● Obtain an IP address automatically

○ Use the following IP address:

IP address: . . .

Subnet mask: . . .

Default gateway: . . .

● Obtain DNS server address automatically

○ Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

□ Validate settings upon exit

Figure 11.1-8: Obtain an IP address automatically

14. Click "**OK**" button.
15. Still in Lab NIC window, click the "**Details…**" button.

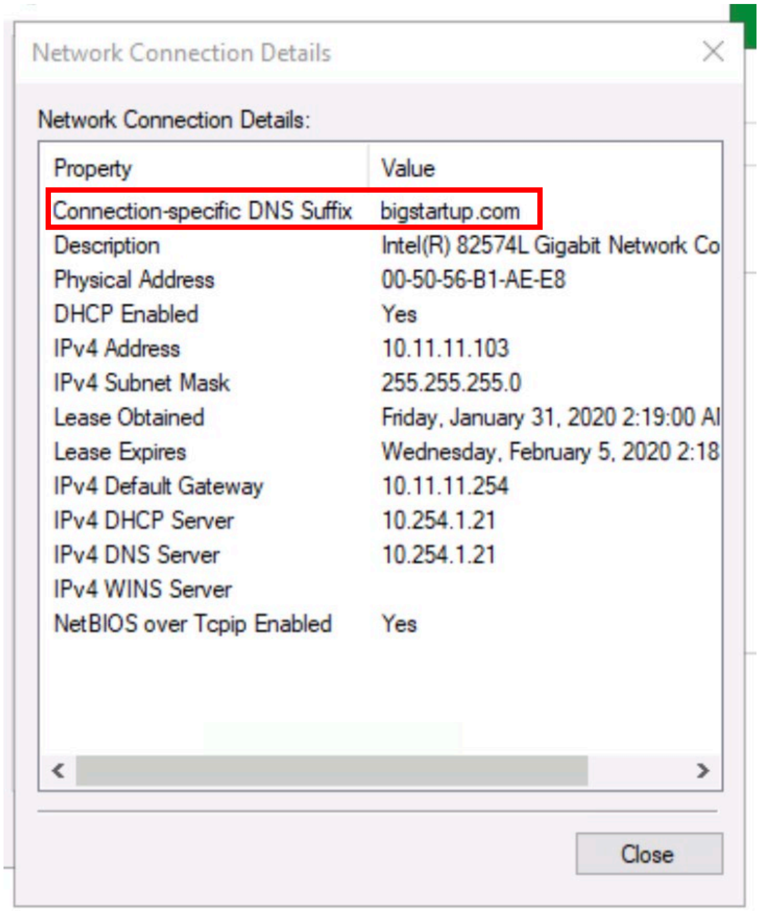

Network Connection Details:

Property	Value
Connection-specific DNS Suffix	bigstartup.com
Description	Intel(R) 82574L Gigabit Network Co
Physical Address	00-50-56-B1-AE-E8
DHCP Enabled	Yes
IPv4 Address	10.11.11.103
IPv4 Subnet Mask	255.255.255.0
Lease Obtained	Friday, January 31, 2020 2:19:00 AI
Lease Expires	Wednesday, February 5, 2020 2:18
IPv4 Default Gateway	10.11.11.254
IPv4 DHCP Server	10.254.1.21
IPv4 DNS Server	10.254.1.21
IPv4 WINS Server	
NetBIOS over Tcpip Enabled	Yes

Figure 11.1-9: Network Connection Details

What connection-specific DNS Subfix did you get?

____

What IP address and subnet mask did you get?

____

What IPv4 Server was assigned to PC-3?

16. Click “**Close”**.

PC-4

17. Access PC-4.
18. Repeat steps 5 to 17.

Task 5: Save Your Configurations

Objectives

You will now proceed to save your configuration.

Steps

Core-1 and Core-2 (via PC-1)

1. Save the current Cores' configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

You have completed Lab 11.1!

AOS-CX Switching Fundamentals

Lab 11.2: OSPF Advanced Settings (Optional)

Overview

When Rent4Cheap Properties started connecting several customers through the same core switches, and each decided to run their own suite of features and protocols (VRFs, VRRP, OSPF, BGP, ACLs, QoS, DC functions, etc). System resources (CPU and memory) rose twenty percent. Although still under normal thresholds, Rent4Cheap Properties would like their customers to make their configurations more resource efficient.

BigStartup also wants to enable symmetric routing, deterministic load balancing on links to the server switch and eliminate the risk of Layer 3 loops for traffic to the internet. Remember, dual floating routes for internet access deployed on Core switches in Lab 9.2 do offer redundancy but also introduce the chance of loops if both ISPs go down. Management also wants to prevent users from receiving OSPF related packets.

You have been asked to optimize the configuration on the Core switch pair. Locally, all these changes are OSPF related.

Objectives

After completing this lab, you will be able to:

- Manipulate paths
- Create loopback interfaces
- Enable passive interfaces
- Change network type to point to point
- Inject a default prefix through OSPF

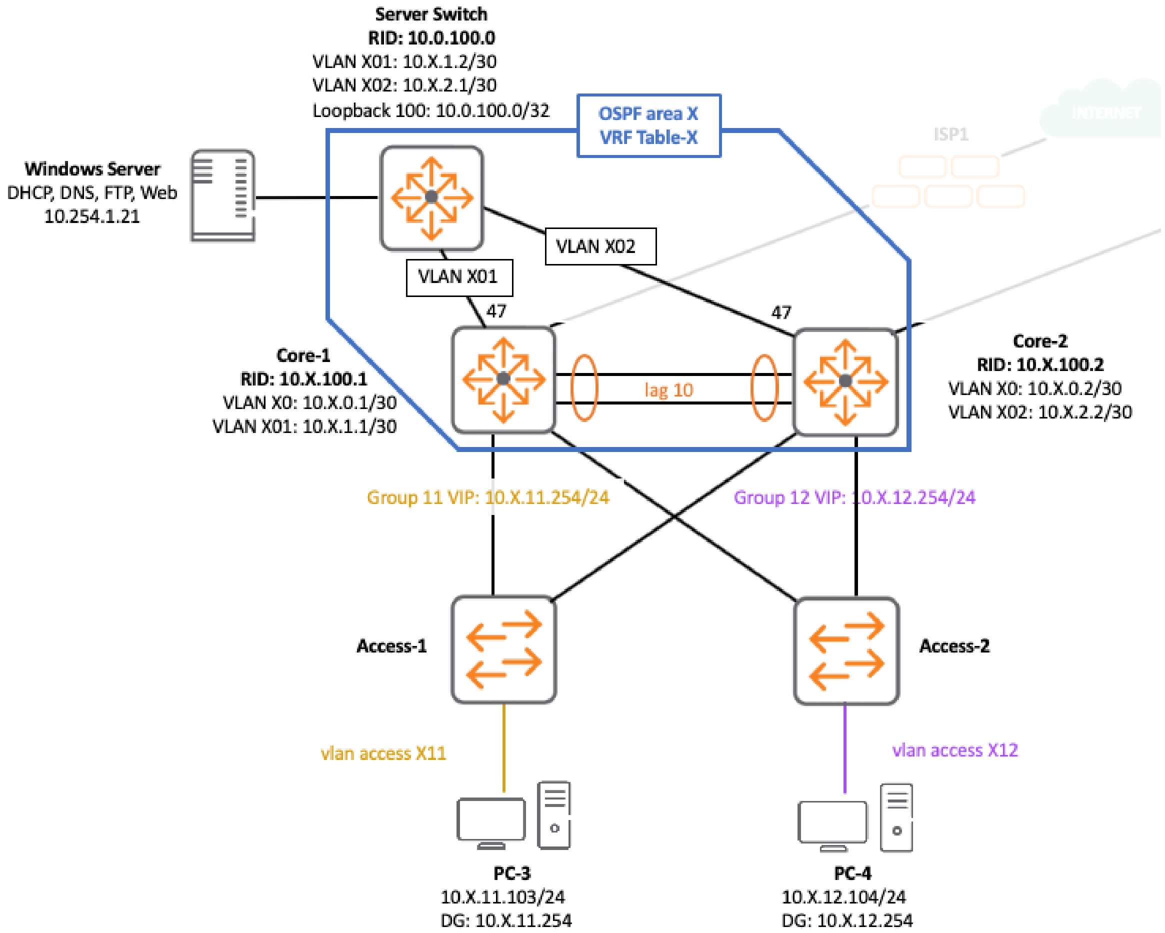

Figure 11.2-1: Lab Topology.

Task 1: Cost based Path Manipulation (traffic engineering).

Objectives

In this activity you will analyze the OSPF link costs and routing table and validate the traffic paths to and from the servers.

BigStartup says response time varies over VLANs X11 and X12 when accessing the servers to the point where applications disconnect. However, when testing the service locally within the same Server segment it works flawlessly. They have reported the issue to Rent4Cheap Properties, and after some research, they confirm there is high bandwidth utilization on one of the two links.

You run some tests and get the following results from PC-3 and PC-4 respectively.

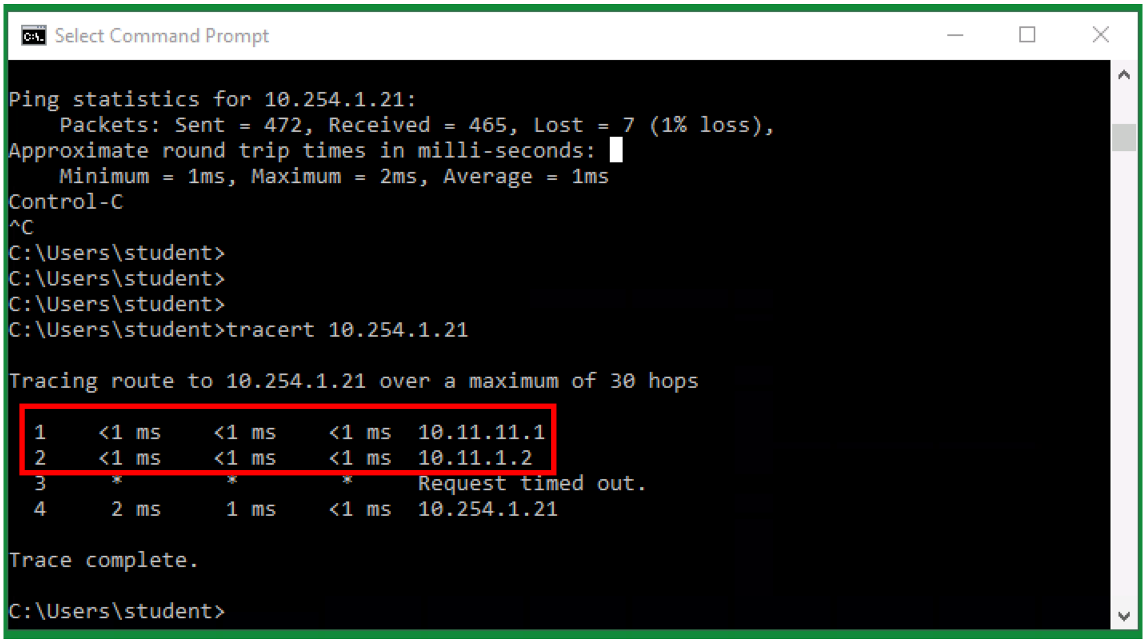

```
Ping statistics for 10.254.1.21:
    Packets: Sent = 472, Received = 465, Lost = 7 (1% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 2ms, Average = 1ms
Control-C
^C
C:\Users\student>
C:\Users\student>
C:\Users\student>
C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.1.2
  3     *        *        *     Request timed out.
  4     2 ms     1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>
```

Figure 11.2-2: PC-3's traceroute

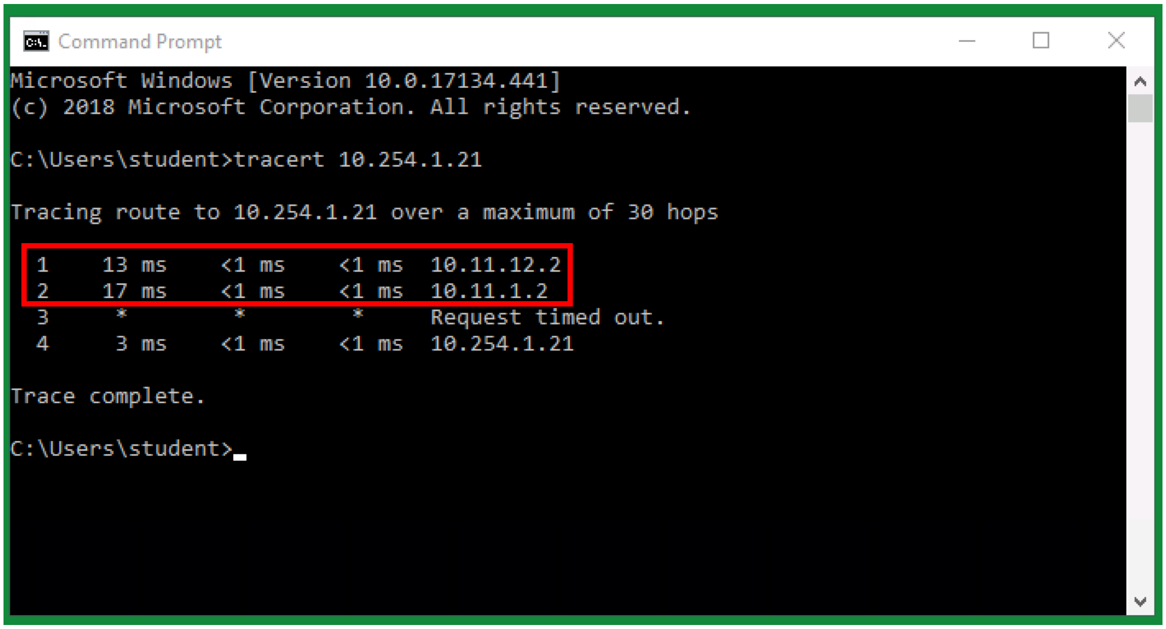

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    13 ms    <1 ms    <1 ms  10.11.12.2
  2    17 ms    <1 ms    <1 ms  10.11.1.2
  3     *        *        *     Request timed out.
  4     3 ms    <1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>
```

- **Figure 11.2-3: PC-4's traceroute**

Figure 11.2.2 shows the expected output, traffic goes from PC-3 to Core-1 on VLAN X11, to Server Switch on VLAN X01, to the Server.

However, figure 11.2.3 shows traffic going from PC-4 to Core-2 on VLAN X11, to Server Switch on VLAN X01, to the Server. The second hop looks strange because Core-2 is not on VLAN X01 at all.

In order to understand traffic paths, it is always better to look at the topology and compare path costs.

Steps

Core-1 (via PC-1)

1. Open an SSH session of Core-1.
2. Display the OSPF interfaces' details in your VRF. Focus on interfaces costs.

```
Core-1# show ip ospf X interface vrf TABLE-X | begin 6 Interface
Interface vlan110 is up, line protocol is up
---------------------------------------------

IP address 10.11.0.1/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1101 is up, line protocol is up
---------------------------------------------

IP address 10.11.1.1/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1111 is up, line protocol is up
----------------------------------------------

IP address 10.11.11.1/24, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Backup-dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1112 is up, line protocol is up
----------------------------------------------

IP address 10.11.12.1/24, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Backup-dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Core-1#
```

3. Write down costs for these four interfaces in figure 11.2-4.

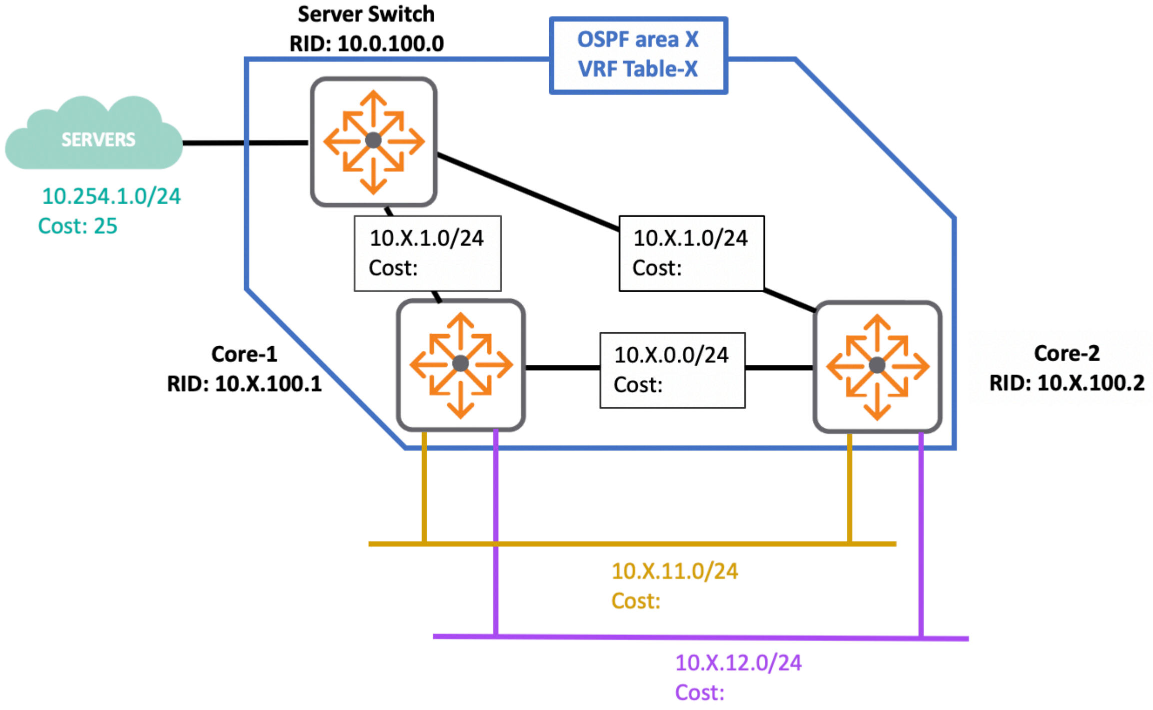

- Figure 11.2-4: Link costs

Core-2 (via PC-2)

4. Open an SSH session of Core-2.
5. Display your VRF OSPF interface details. Focus on interface costs.

```
Core-2# show ip ospf X interface vrf TABLE-X | begin 6 Interface
Interface vlan110 is up, line protocol is up
---------------------------------------------

IP address 10.11.0.2/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1111 is up, line protocol is up
----------------------------------------------

IP address 10.11.11.2/24, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1112 is up, line protocol is up
----------------------------------------------
```

```
IP address 10.11.12.2/24, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Interface vlan1102 is up, line protocol is up
---------------------------------------------

IP address 10.11.2.2/30, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
Core-2#
```

6. Confirm Core-1 and Core-2 use the same link costs they have in common (**VLAN X0, X11 and X12**). Then write down in figure 11.2-4 the link cost to Server Switch (**10.X.2.0/30**).

NOTE: The only missing link cost is the Server segment. However, you were told that cost is 25.

This information can be used to predict traffic paths. For Core-2 there are two options for reaching the servers, the path via Core-1 with a total cost of 225 (100+100+25) or the path through Server Switch with a total cost of 125 (100+25). When running OSPF if there are two paths of the same type (intra Area OSPF in this case), the one with lowest cost is preferred and published in both the OSPF routing table and also in the VRF (or global) routing table. Therefore, Core-2 uses the server switch in VLAN X02 (10.X.2.1) as its next hop (figure 11.2-5). You saw this in Lab 11.1 Task 2 Step 21:

```
Core-2# show ip route ospf vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.0.100.0/32, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf
10.11.1.0/30, vrf TABLE-11
        via  10.11.0.1,  [110/200],  ospf
        via  10.11.2.1,  [110/200],  ospf
10.254.1.0/24, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf

Core-2#
```

Likewise, Core-1 will use the path with the lowest cost. This is through VLAN X01 (figure 11.2-6).

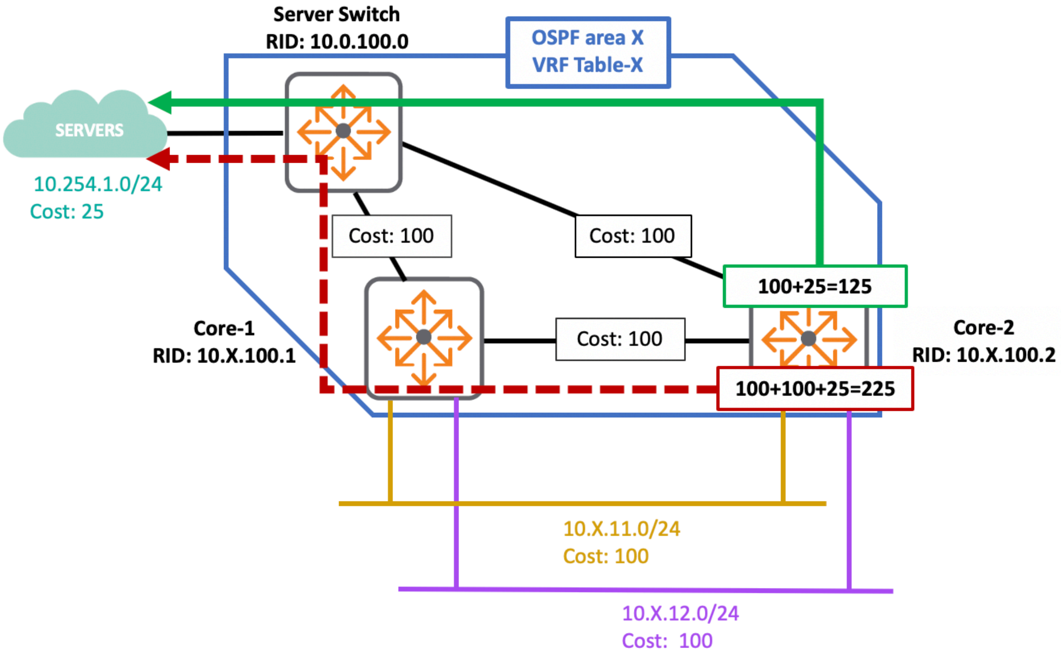

- Figure 11.2-5: Path via VLAN X02

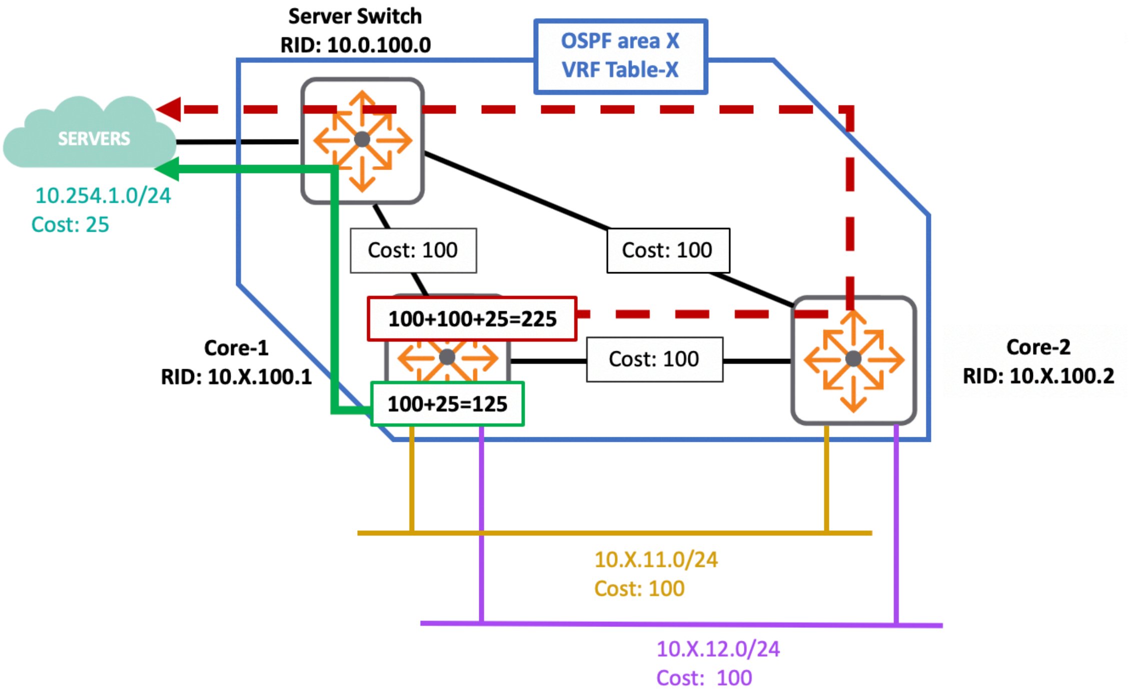

- **Figure 11.2-6: Path via VLAN X01**

NOTE: The Server Switch on the other hand has two options for reaching VLANs X11 and X12. It can use Core-1 or Core-2. Each has a total cost of 200. You can inspect the routing table to validate this.

7. Confirm Core-1 and Core-2 use the same costs for links they have in common (VLAN X0, X11 and X12). Then write down in figure 11.2-4 the link cost to the Server Switch (**10.X.2.0/30**).

Server Switch (via PC-1)

8. Using Putty, open an SSH session to **Server Switch.**

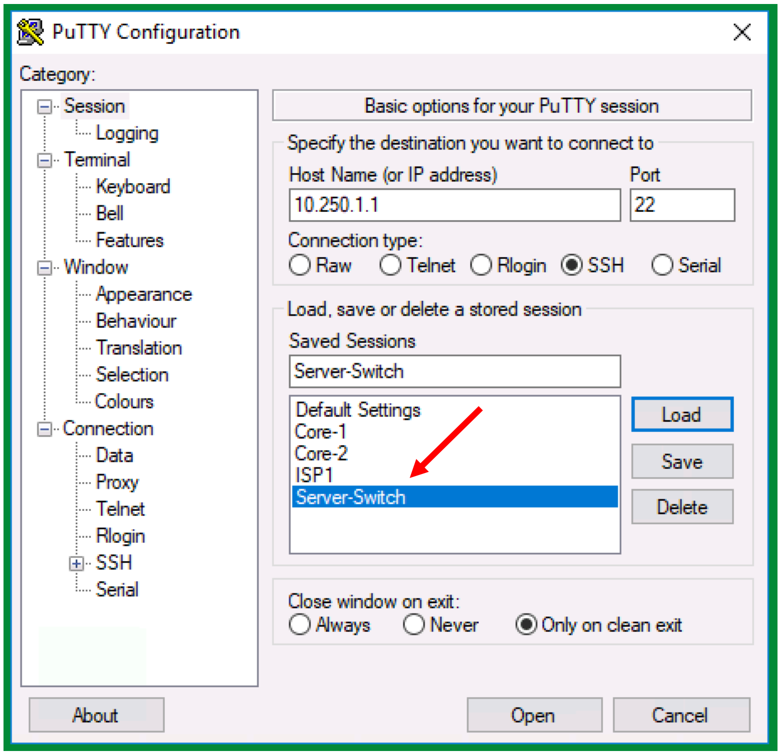

Host Name (or IP address): 10.250.1.1
Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Saved Sessions: Server-Switch

Default Settings
Core-1
Core-2
ISP1
Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 11.2-7: SSH to Server Switch

9. Login using **cxfX/aruba123.**

> **NOTE:** Replace the highlighted “X” for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

10. Display the routing table of vrf **CXF_SERVER-SWITCH_TABLE-X**, including specifically segments **10.X.11.0/24 and 10.X.12.0/24**.

```
P52-6300-AB# show ip route 10.X.11.0/24 vrf CXF_SERVER-SWITCH_TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.11.0/24, vrf CXF_SERVER-SWITCH_TABLE-11
        via  10.11.1.1,  [110/200],  ospf
        via  10.11.2.2,  [110/200],  ospf

P52-6300-AB#
```

```
P52-6300-AB# show ip route 10.X.12.0/24 vrf CXF_SERVER-SWITCH_TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.12.0/24, vrf CXF_SERVER-SWITCH_TABLE-11
        via  10.11.1.1,  [110/200],  ospf
        via  10.11.2.2,  [110/200],  ospf

P52-6300-AB#
```

How many next hops do the 10.11.11.0/24 and 10.11.12.0/24 networks have?

What is the cost in both cases?

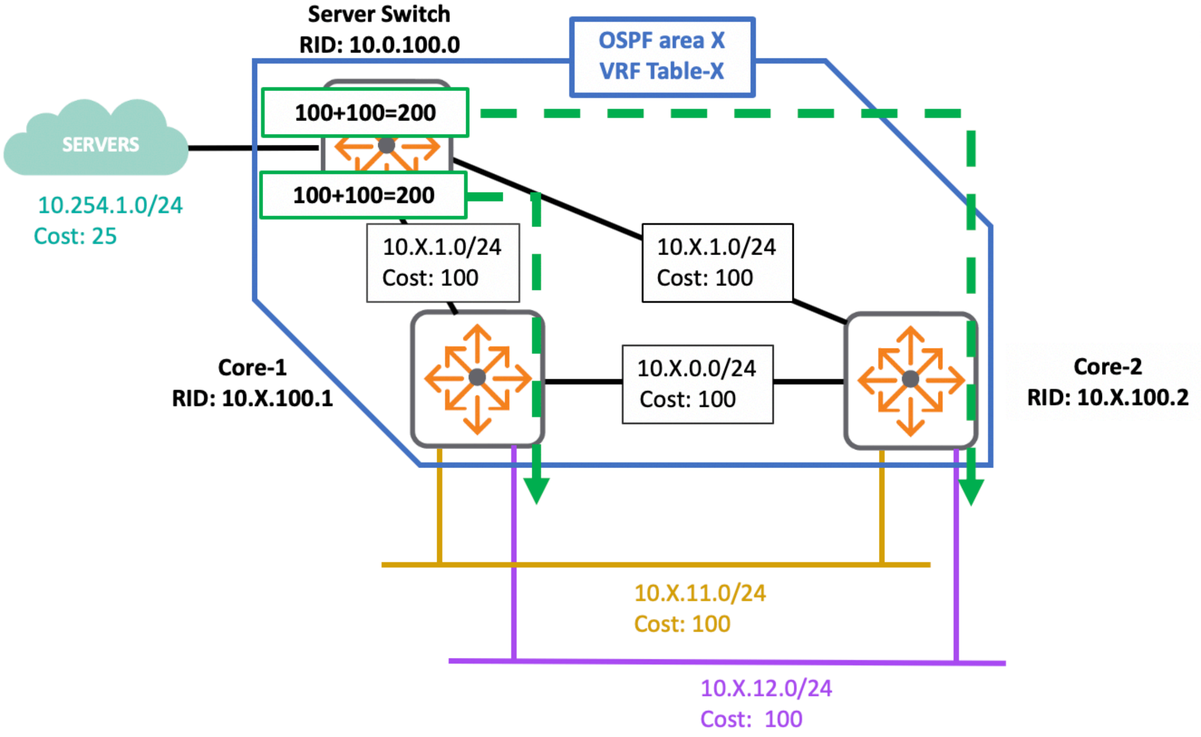

Figure 11.2-8: ECMP

NOTE: The server switch has two alternatives for sending traffic to the LAN segments and will perform ECMP to balance the load using a flow-based algorithm. Unfortunately, you do not have control of those decisions, and it might lead to asymmetric multi-hop routing, which in turn can generate delay and jitter. Also, if a firewall appliance is to be deployed between the LAN segments and the servers, then it is important that connection flows always use the same interface inbound and outbound, otherwise the firewall could drop valid traffic.

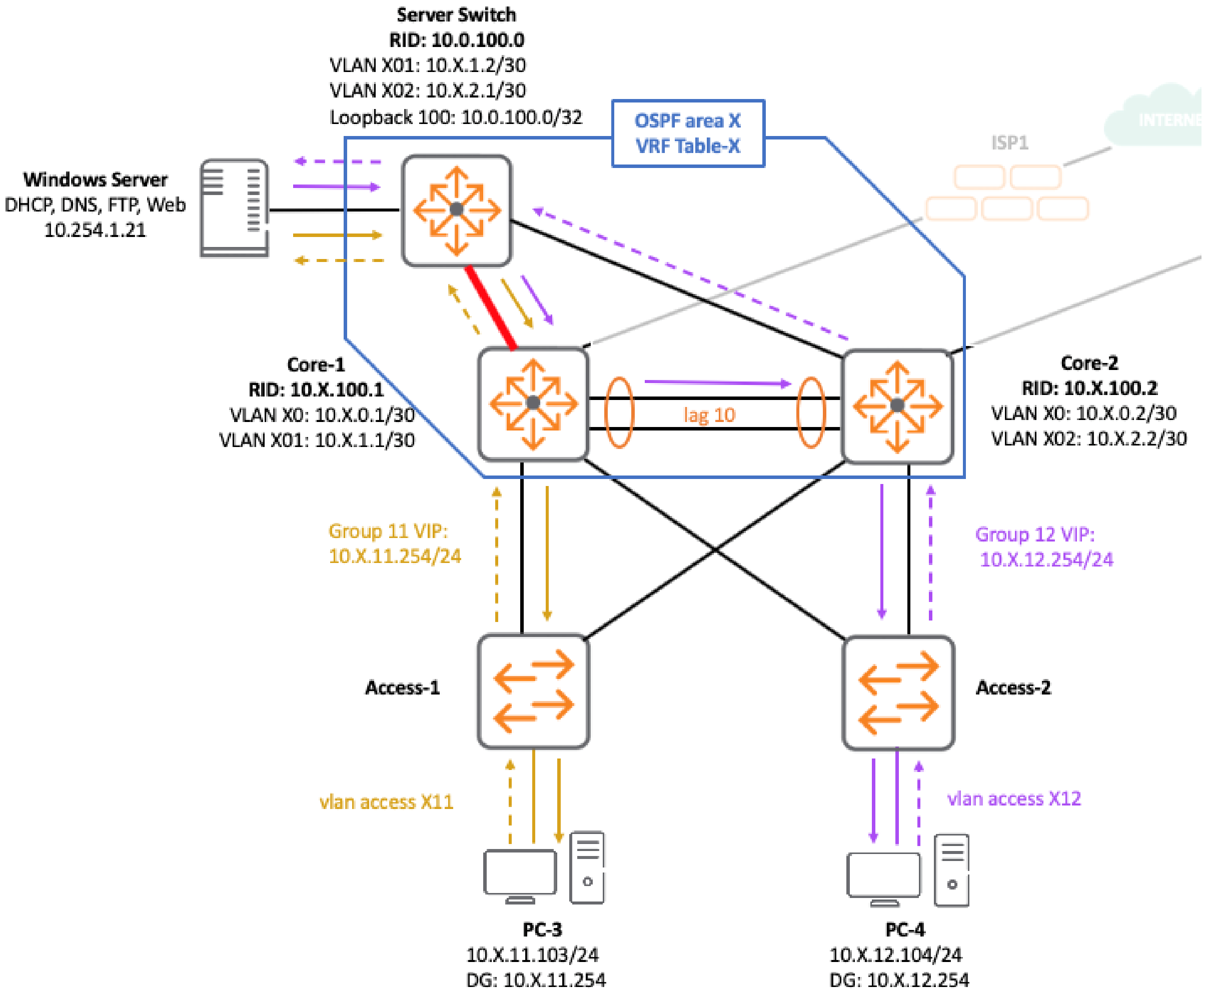

Figure 11.2-9: Asymmetric multi-hop routing

NOTE: Your customer desires more control over what path the traffic is using. As the administrator, you can influence routing decisions by manually changing the costs and making some paths more preferred.

To make traffic use VLAN X11 on Core-1, reduce the cost Core-1 advertises for the VLAN X11 link in its Router LSA. This will make the Server Switch calculate a lower overall path cost through Core-1 vs Core-2.

To make traffic use VLAN X12 on Core-2, reduce the cost Core-2 advertises for the VLAN X12 link in its Router LSA. This will make the Server Switch calculate a lower overall path cost through Core-2 vs Core-1.

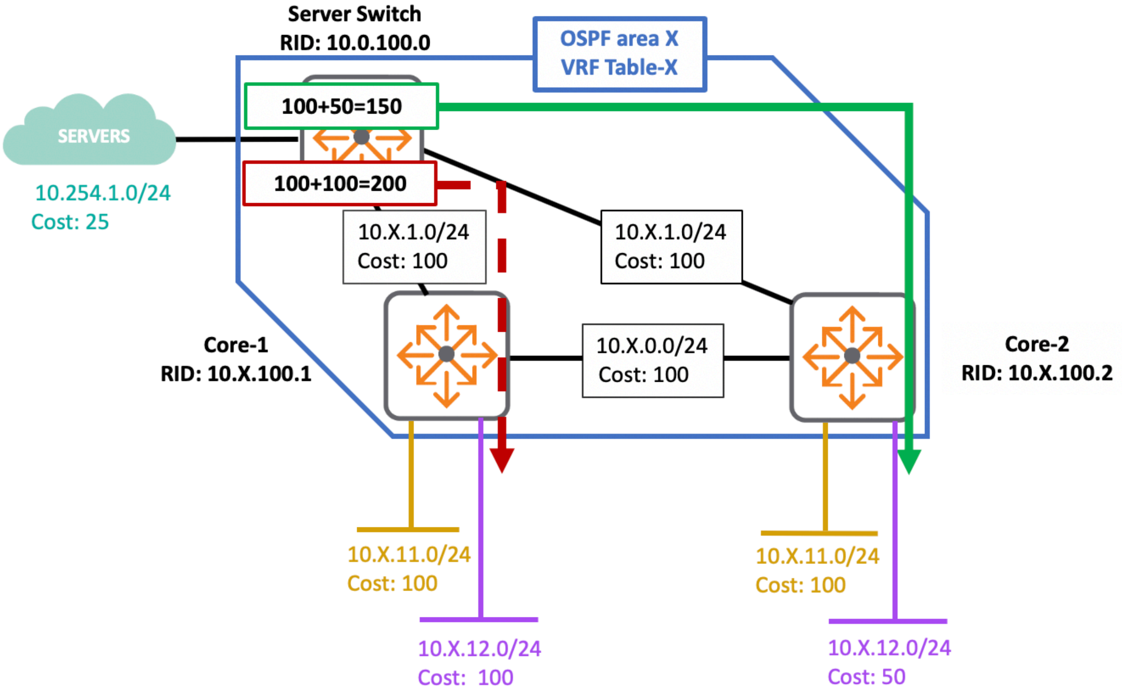

- **Figure 11.2-10: Path via Core-2**

Core-2 (via PC-1)

11. Move **to Core-2.**
12. Reduce the OSPF cost of interface **VLAN X12** to 50.

```
Core-2# configure terminal
Core-2(config)# interface vlan X12
Core-2(config-if-vlan)# ip ospf cost 50
Core-2(config-if-vlan)# end
Core-2#
```

13. Use the **show ip ospf interface** command for validating the change. Notice how the output says the new value was configured.

```
Core-2# show ip ospf interface vlanX12 vrf TABLE-X
Interface vlan1112 is up, line protocol is up
----------------------------------------------

IP address 10.11.12.2/24, Process ID 11 VRF TABLE-11, area 0.0.0.11
    State Dr-other, Status up, Network type Broadcast
    Link Speed: 1000 Mbps
    Cost Configured 50, Calculated 50
    Transit delay 1 sec, Router priority 1
    No designated router on this network
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled

Core-2#
```

Core-1 (via PC-1)

14. Move back to Core-1.
15. Reduce the OSPF cost of interface **VLAN X11** to 50.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11
Core-1(config-if-vlan)# ip ospf cost 50
Core-1(config-if-vlan)# exit
```

Server Switch (via PC-1)

16. Move back to the Server Switch.
17. Display the routing table of vrf **CXF_SERVER-SWITCH_TABLE-X**, specifically including segments **10.X.11.0/24** and **10.X.12.0/24**.

```
P52-6300-AB# show ip route 10.X.11.0/24 vrf CXF_SERVER-SWITCH_TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.11.0/24, vrf CXF_SERVER-SWITCH_TABLE-11
        via  10.11.1.1,  [110/150],  ospf

P52-6300-AB#
```

```
P52-6300-AB# show ip route 10.X.12.0/24 vrf CXF_SERVER-SWITCH_TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.12.0/24, vrf CXF_SERVER-SWITCH_TABLE-11
        via  10.11.2.2,  [110/150],  ospf

P52-6300-AB#
```

How many next hops do the 10.11.11.0/24 and 10.11.12.0/24 networks have?

What is the total cost to those prefixes?

PC-3 and PC-4

18. Access PC-3.
19. Run a traceroute to the server **10.254.1.21**.

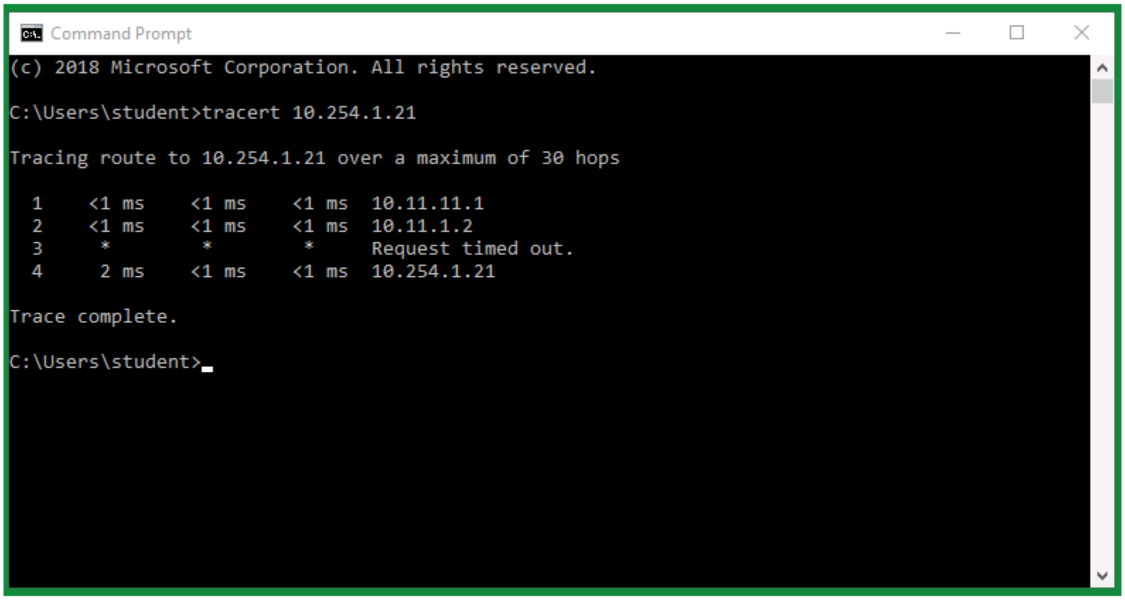

```
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.1.2
  3     *        *        *     Request timed out.
  4     2 ms    <1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>_
```

- **Figure 11.2-11: Tracert**

20. Access PC-4.
21. Run a traceroute to the **server 10.254.1.21**.

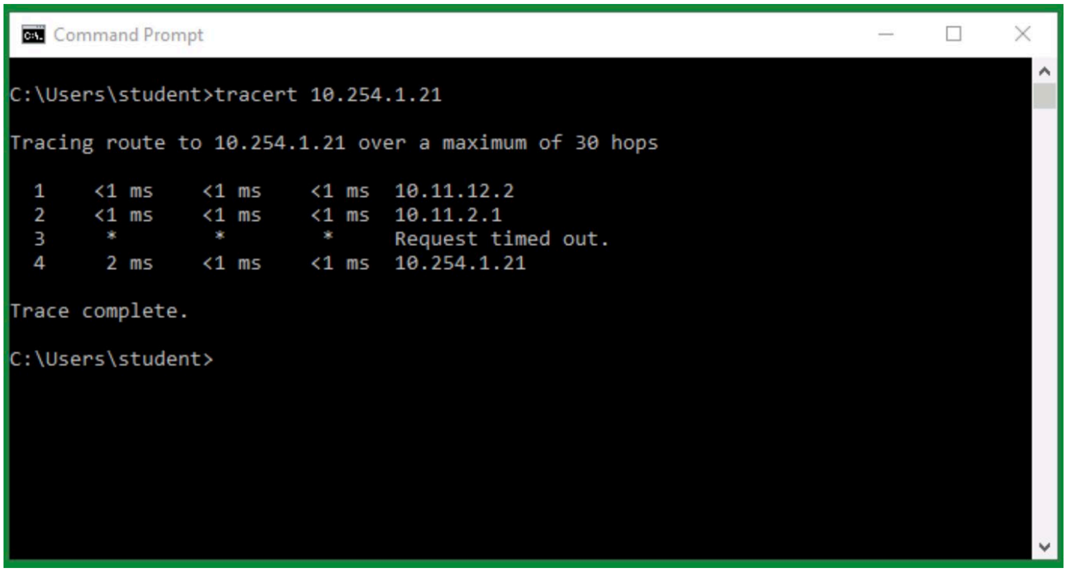

```
C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.12.2
  2    <1 ms    <1 ms    <1 ms  10.11.2.1
  3     *        *        *     Request timed out.
  4     2 ms    <1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>
```

- **Figure 11.2-12: Tracert**

What are the first and second hop of PC-3?

____

What are the first and second hop of PC-4?

Are these results you expected?

Next, test redundancy by disabling the Layer 3 connection between Core-1 and the Server Switch.

Core-1 (via PC-1)

22. Move back to Core-1.
23. Disable interface **VLAN X01**.

```
Core-1(config)# interface vlan X01
Core-1(config-if-vlan)# shutdown
```

TIP: Since the physical port between both devices (1/1/47) remains up, the Server Switch hasn’t sensed the failure yet. You will have to wait 40 seconds before moving forward, which is the value of the Dead timer.

In production scenarios you would normally rely on Bidirectional Forwarding Detection (BFD) in order to detect down neighbors regardless of the state of the physical media. BFD is covered in the Implementing Aruba Switching training course.

PC-3

24. Move back to PC-3.
25. Run a traceroute to the server (**10.254.1.21)**. Test should be successful.

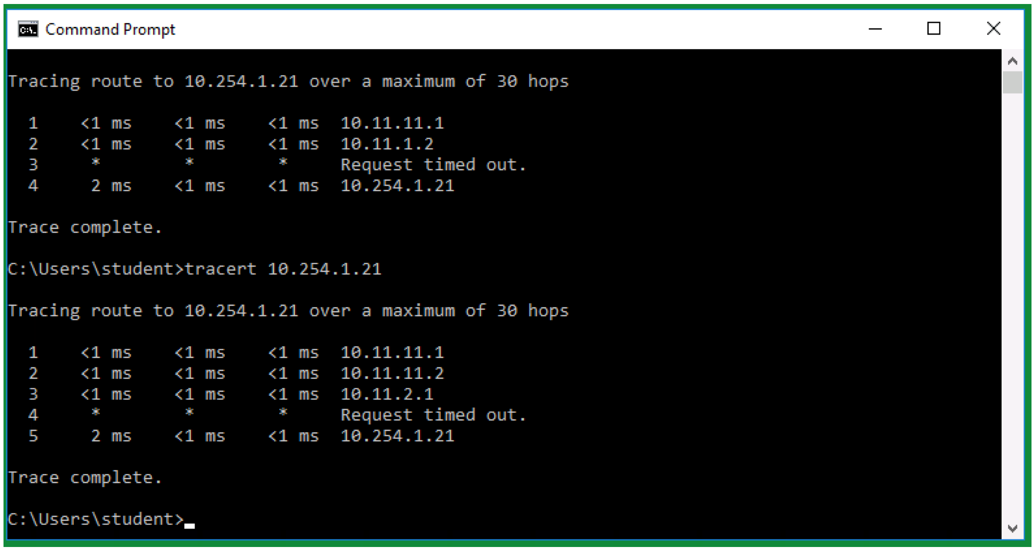

```
Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.1.2
  3     *        *        *     Request timed out.
  4     2 ms    <1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>tracert 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.11.11.1
  2    <1 ms    <1 ms    <1 ms  10.11.11.2
  3    <1 ms    <1 ms    <1 ms  10.11.2.1
  4     *        *        *     Request timed out.
  5     2 ms    <1 ms    <1 ms  10.254.1.21

Trace complete.

C:\Users\student>_
```

- Figure 11.2-13: Tracert

What are the first three hops now?

____

Is this result you expected?

____

Server Switch (via PC-1)

26. Move back to the Server Switch.
27. Display the routing table of vrf **CXF_SERVER-SWITCH_TABLE-X**, including specifically the segment **10.X.11.0/24**.

```
P52-6300-AB# show ip route 10.X.11.0/24 vrf CXF_SERVER-SWITCH_TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.11.0/24, vrf CXF_SERVER-SWITCH_TABLE-11
        via  10.11.2.2,  [110/200],  ospf

P52-6300-AB#
```

What is the next-hop now?

What is the total cost to that prefix?

Core-1 (via PC-1)

28. Move back to Core-1.
29. Enable interface **VLAN X01**.

```
Core-1(config-if-vlan)# no shutdown
Core-1(config-if-vlan)# end
Core-1#
```

Task 2: Passive Interfaces

Objectives

When enabling OSPF on a Layer 3 interface, there will be two immediate results: first the link's segment gets included in the Router LSA. Second the router will start advertising hello packets periodically based in the hello interval configured for that interface.

However, there are links where sending those messages is not necessary and can even introduce security risks.

That is the case of the LAN segments where hosts reside. Since hello messages use local link scoped multicast packets for both hello and Link State Updates, any host will receive those messages when they are sent out the on the VLAN. If somebody is running packet analysis software, they could see the contents and perform a reconnaissance attack, a Denial of Service (DoS) attack, or man in the middle attack.

By suppressing hello messages on VLANs X11 and X12 you will improve security as well as control plane and data plane performance. Data plane performance is improved by preventing the segments from being considered "transit" networks.

Steps

PC-3

1. Open a console session to PC-3.
2. Open **Wireshark**, there should be a shortcut on the Desktop.
3. Double click the "**Lab NIC**" entry. That will begin the packet capture on that interface.

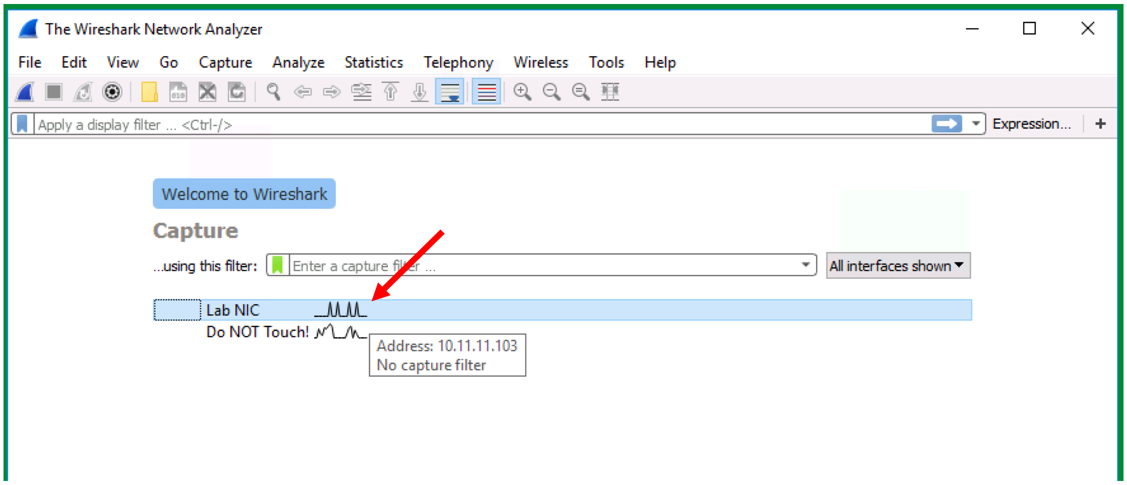
The Wireshark Network Analyzer
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Apply a display filter ... <Ctrl-/>
Expression...
Welcome to Wireshark
Capture
...using this filter:
Enter a capture filter ...
All interfaces shown
Lab NIC
Do NOT Touch!
Address: 10.11.11.103
No capture filter

Figure 11.2-14: Wireshark.

4. On filter type “**ospf**” with no quotes and hit **[Enter]**. That will instruct Wireshark to only present OSPF packets. Wait a few seconds and you will start to see hello packets every 10 seconds.

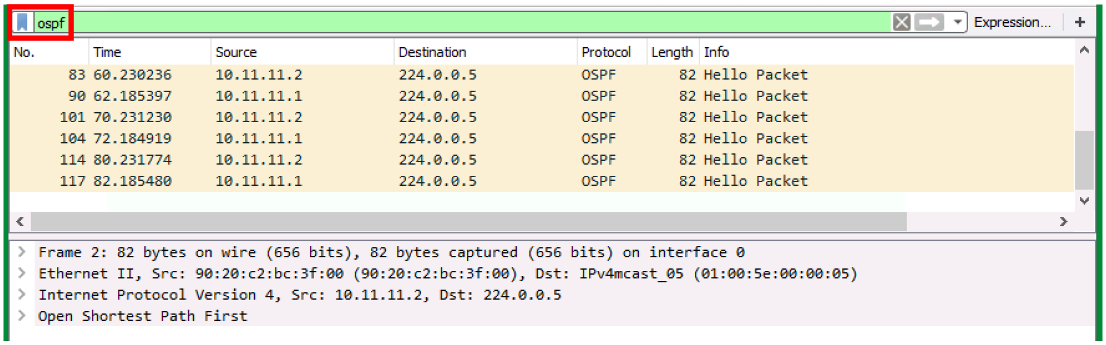

```
No.   Time        Source        Destination   Protocol  Length  Info
   83 60.230236   10.11.11.2    224.0.0.5     OSPF          82  Hello Packet
   90 62.185397   10.11.11.1    224.0.0.5     OSPF          82  Hello Packet
  101 70.231230   10.11.11.2    224.0.0.5     OSPF          82  Hello Packet
  104 72.184919   10.11.11.1    224.0.0.5     OSPF          82  Hello Packet
  114 80.231774   10.11.11.2    224.0.0.5     OSPF          82  Hello Packet
  117 82.185480   10.11.11.1    224.0.0.5     OSPF          82  Hello Packet
```

```
> Frame 2: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: 90:20:c2:bc:3f:00 (90:20:c2:bc:3f:00), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
> Internet Protocol Version 4, Src: 10.11.11.2, Dst: 224.0.0.5
> Open Shortest Path First
```

Figure 11.2-15: Hello packets.

TIP: Although not the main goal of this task, you can leverage our captures and analyze one of the hello packets for academic purposes.

5. Stop the capture, select one of the packets and expand the transport header row (Open Shortest Path First), then OSPF Hello Packet underneath.

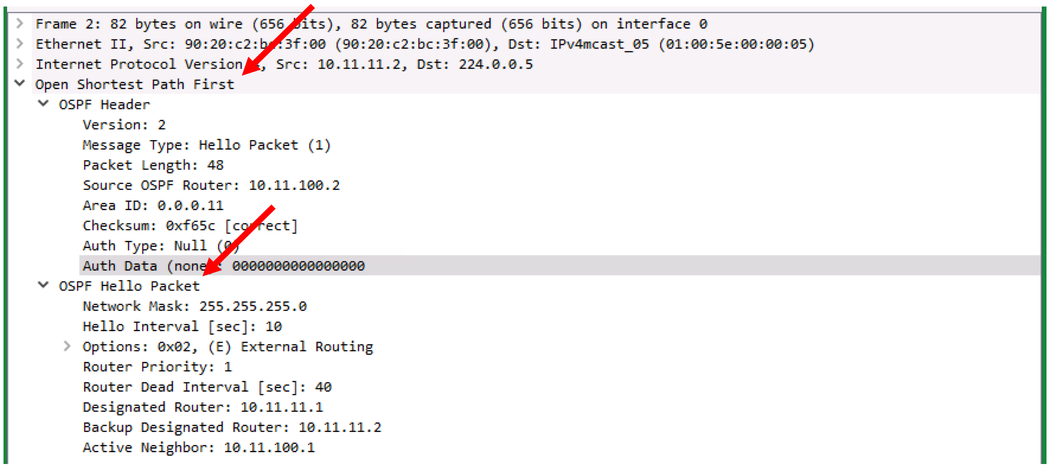

```
> Frame 2: 82 bytes on wire (656 bits), 82 bytes captured (656 bits) on interface 0
> Ethernet II, Src: 90:20:c2:bc:3f:00 (90:20:c2:bc:3f:00), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
> Internet Protocol Version 4, Src: 10.11.11.2, Dst: 224.0.0.5
v Open Shortest Path First
  v OSPF Header
      Version: 2
      Message Type: Hello Packet (1)
      Packet Length: 48
      Source OSPF Router: 10.11.100.2
      Area ID: 0.0.0.11
      Checksum: 0xf65c [correct]
      Auth Type: Null (0)
      Auth Data (none): 0000000000000000
  v OSPF Hello Packet
      Network Mask: 255.255.255.0
      Hello Interval [sec]: 10
    > Options: 0x02, (E) External Routing
      Router Priority: 1
      Router Dead Interval [sec]: 40
      Designated Router: 10.11.11.1
      Backup Designated Router: 10.11.11.2
      Active Neighbor: 10.11.100.1
```

Figure 11.2-16: Hello packet's contents.

What protocol version it is?

What packet type are you watching?

What is the Source address?

What is the Area ID?

What are the authentication type and data?

What is the Network Mask, and why is it included?

____

What is the Dead interval?

____

What is included in the Neighbor list?

____

IMPORTANT: There are a few attributes within the hello messages that are critical for successfully establishing neighbor relationships.

a) Attributes that must be different: Router ID
b) Attributes that must be identical: Version, Area #, Authentication type and data, Area flags, subnet mask, Hello and Dead intervals.
c) Attributes that can be the same or different: Priority, Designated and Backup Designated routers and Neighbor List.

When a neighbor relationship is not coming up between two OSPF routers that reside within the same segment (as in the case of lab 10.1 task 2), step back and check these values before looking at anything else.

Core-1 (via PC-1)

6. Move to Core-1.
7. Look at the Link State Database process in your VRF.

```
Core-1# show ip ospf lsdb vrf TABLE-X area X
OSPF Router with ID (10.11.100.1) (Process ID 11 VRF TABLE-11)
=============================================================
Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.0.100.0      10.0.100.0      285       0x80000043 0x00007486     2
10.11.100.1     10.11.100.1     21        0x8000007d 0x00000dd3     4
10.11.100.2     10.11.100.2     17        0x80000086 0x00005386     4

Network Link State Advertisements (Area 0.0.0.11)
--------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
--------------------------------------------------------------
10.11.0.1       10.11.100.1     1437      0x8000005c 0x0000379d
10.11.1.1       10.11.100.1     289       0x8000001b 0x0000e706
10.11.2.2       10.11.100.2     412       0x8000003c 0x0000b429
10.11.11.2      10.11.100.2     22        0x80000001 0x0000d1f5
10.11.12.2      10.11.100.2     22        0x80000001 0x0000c8fe

Core-1#
```

How many Router LSAs do you count?

____

How many Network LSAs do you count?

____

8. Use the **show ip ospf neighbors**.

```
Core-1# show ip ospf neighbors vrf TABLE-X
OSPF Process ID 11 VRF TABLE-11
================================

Total Number of Neighbors: 4

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.11.100.2      1         FULL/BDR          10.11.0.2          vlan110

10.0.100.0       1         FULL/BDR          10.11.1.2          vlan1101

10.11.100.2      1         FULL/BDR          10.11.11.2         vlan1111
```

```
10.11.100.2       1         FULL/BDR          10.11.12.2         vlan1112

Core-1#
```

How many neighbors does Core-1 have?

9. Set the SVIs of **VLAN X11** and **X12** passive interfaces.

```
Core-1# configure terminal
Core-1(config)# interface vlan X11-X12
Core-1(config-if-vlan-<1111-1112>)# ip ospf passive
Core-1(config-if-vlan-<1111-1112>)# end
```

10. Display the neighbor list on your VRF again.

```
Core-1# show ip ospf neighbors vrf TABLE-X

OSPF Process ID 11 VRF TABLE-11
================================

Total Number of Neighbors: 2

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.11.100.2       1         FULL/BDR          10.11.0.2          vlan110

10.0.100.0        1         FULL/BDR          10.11.1.2          vlan1101

Core-1#
```

How many neighbors does Core-1 have now?

Core-2 (via PC-1)

11. Move to Core-2.

12. Repeat step 12 using **VLANs X11** and **X12.**

```
Core-2# configure terminal
Core-2(config)# interface vlan X11-X12
Core-2(config-if-vlan-<1111-1112>)# ip ospf passive
Core-2(config-if-vlan-<1111-1112>)#
```

PC-3

13. Move back to PC-3.
14. Start a new capture, then wait a minute. You will notice there are no more OSPF packets showing up.

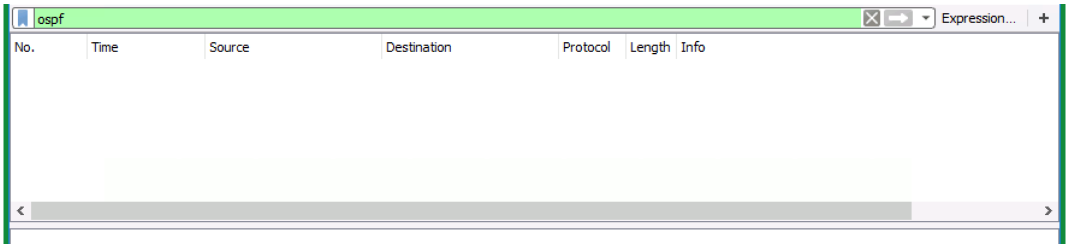

Figure 11.2-17: Wireshark output.

15. Look at the Link State Database again.

```
Core-1# show ip ospf lsdb vrf TABLE-X area X
OSPF Router with ID (10.11.100.1) (Process ID 11 VRF TABLE-11)
===============================================================

Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------

LSID            ADV Router      Age        Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.0.100.0      10.0.100.0      338        0x80000043 0x00007486       2
10.11.100.1     10.11.100.1     13         0x8000007e 0x0000fe24       4
10.11.100.2     10.11.100.2     9          0x80000087 0x0000ec31       4

Network Link State Advertisements (Area 0.0.0.11)
--------------------------------------------------

LSID            ADV Router      Age        Seq#       Checksum
--------------------------------------------------------------
10.11.0.1       10.11.100.1     1490       0x8000005c 0x0000379d
```

```
10.11.1.1       10.11.100.1     342       0x8000001b 0x0000e706
10.11.2.2       10.11.100.2     465       0x8000003c 0x0000b429

Core-1#
```

How many Network LSAs can you count?

____

Why do you have that number?

____

NOTE: You have two fewer LSAs than before, because, as soon as the Core switches stop seeing each other, VLAN X11 and X12's segments shift to stub. From OSPF's topology perspective both networks will be seen as individually connected behind both switches as in figure 11.2-18.

However, at Layer 2, VLAN X11 and X12's broadcast domains are as in figure 11.2-19

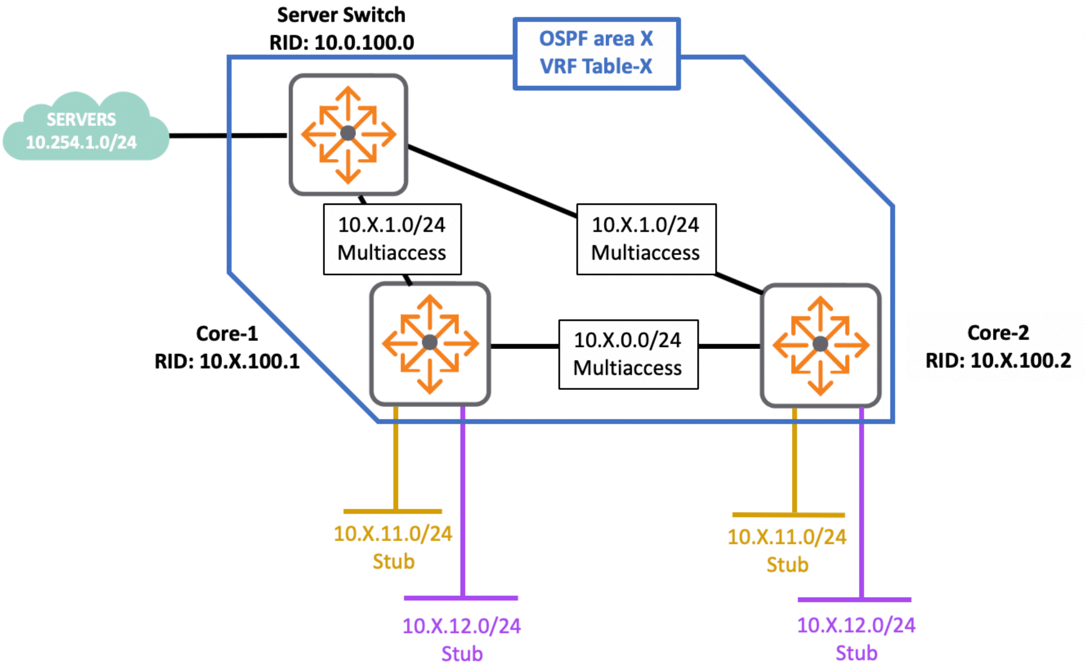

Figure 11.2-18: Stub Networks.

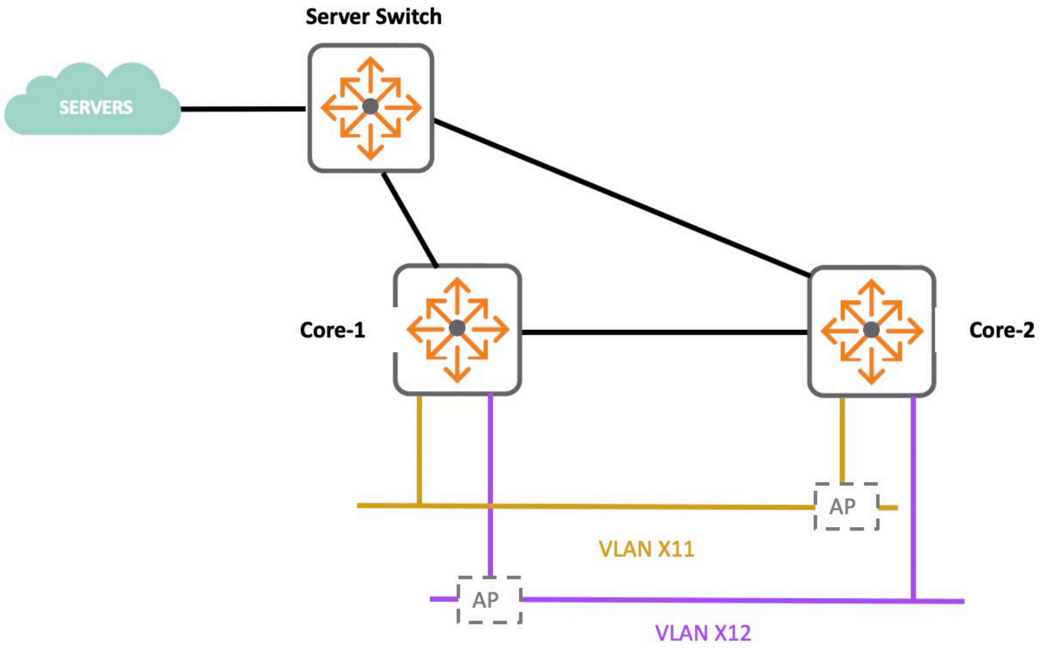

Figure 11.2-19: Layer 2 Topology.

Task 3: Define Point to Point Networks

Objectives

You learned in the OSPF module, when enabling OSPF on routers in a Multiaccess network such as Ethernet or multipoint WAN networks (either physical or Virtual), there will be a DR and BDR election. This election is needed in order to reduce the number of adjacencies within a segment where multiple routers are deployed. However, it is not necessary if there are only two routers.

This process lasts the time defined by the wait interval (usually similar to the dead interval) that covers the amount of time between the links coming up and the DR being elected. That means that no adjacency can happen in that link before this process completes which in turn delays the convergence in critical situations e.g. after a power outage.

Multiaccess Networks have another characteristic, which is their Subnet Masks are not announced within the Router LSA, but in the Network LSA. Therefore, additional LSAs must be created to properly share the topology information, which in turn adds overhead to the overall route selection process when using the Dijkstra algorithm.

If only two routers are present in the segment, DR election is not needed. Nevertheless, it happens because of the type of the network. However, if the administrator knows that no other OSPF devices will be inserted into that broadcast domain, the network type can be changed to point to point.

In point to point networks, as soon as two neighbors discover each other, they begin the LSA exchange immediately and achieve the FULL adjacency state faster. This not only improves convergence time but also makes the routers include the segment's subnet mask in their Router LSA. This eliminates the need for a Network LSA for that link since there will not be a DR to create it.

In this lab you will change VLANs X0, X01 and X02 to point to point.

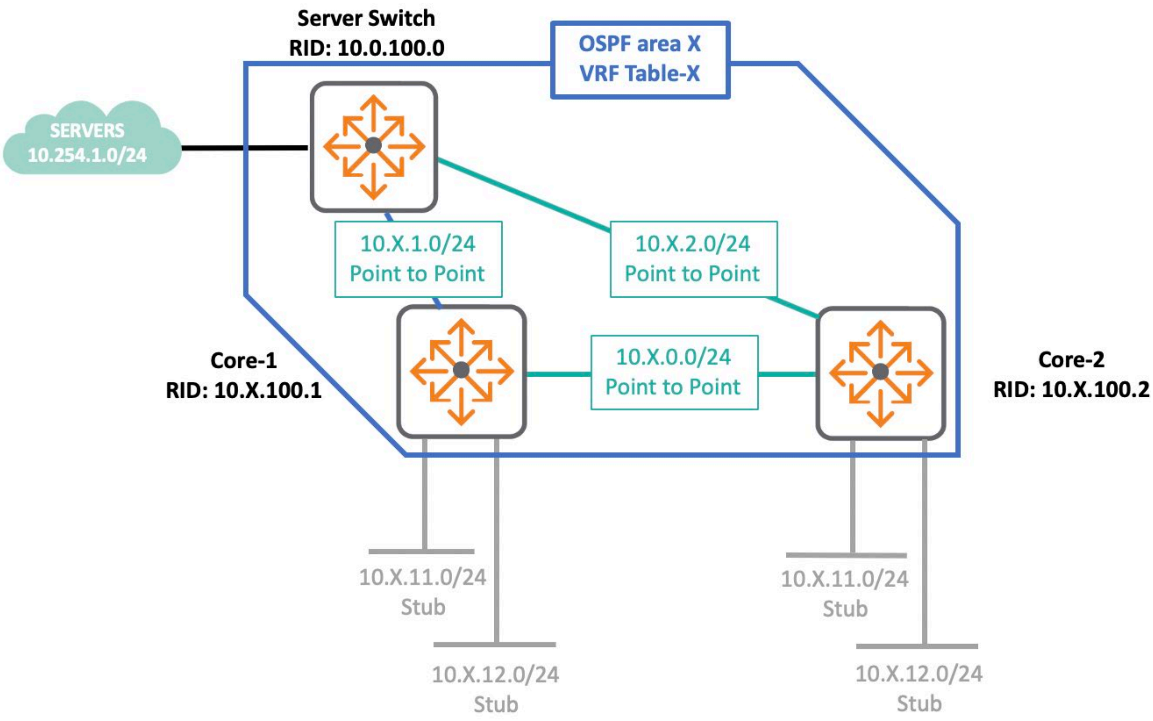

Figure 11.2-20: Point to point links.

Steps

Core-1 (via PC-1)

1. Open a SSH session to Core-1.
2. Change interface **VLAN X0**’s network type to **point-to-point**.

```
Core-1# configure terminal
Core-1(config)# interface vlan X0
Core-1(config-if-vlan)# ip ospf network point-to-point
Core-1(config-if-vlan)#
```

3. Look at the neighbor relationships and focus on the one with the Server Switch.

```
Core-1(config-if-vlan)# do show ip ospf neighbors vrf TABLE-X
OSPF Process ID 11 VRF TABLE-11
================================

Total Number of Neighbors: 2

Neighbor ID      Priority  State              Nbr Address       Interface
```

```
------------------------------------------------------------------------
10.11.100.2       n/a        FULL.             10.11.0.2          vlan110

10.0.100.0        1.         FULL/BDR          10.11.1.2          vlan1101

Core-1(config-if-vlan)#
```

What is Core-1's priority?

____

What is Core-1's role?

____

ANSWER: There is no role. Since the link is point to point, no DR election will happen from Core-1's perspective. Since priority value loses relevance it is omitted from the hello messages.

4. Change interface **VLAN X01**'s network type to **point-to-point**.

```
Core-1(config-if-vlan)# interface vlan X01
Core-1(config-if-vlan)# ip ospf network point-to-point
```

Core-2 (via PC-1)

5. Open an SSH session to Core-2.
6. Change interface **VLAN X0** and **X02**'s network type to **point-to-point**.

```
Core-2(config)# interface vlan X0
Core-2(config-if-vlan)# ip ospf network point-to-point
Core-2(config-if-vlan)# interface vlan X02
Core-2(config-if-vlan)# ip ospf network point-to-point
Core-2(config-if-vlan)# end
```

NOTE: You have completed the Core switch portion of the configuration, still the changes must be made on the Server Switch so all devices in the Area have the same view of the topology and avoid inconsistency. Therefore, you asked the other partner to apply similar commands.

TIP: In next steps you will pretend to be the Server Switch technician.

7. Look at the neighbor relationships.

```
Core-2# show ip ospf neighbors vrf TABLE-X
OSPF Process ID 11 VRF TABLE-11
================================

Total Number of Neighbors: 2

Neighbor ID      Priority  State              Nbr Address       Interface
-------------------------------------------------------------------------
10.11.100.1      n/a       FULL               10.11.0.1          vlan110

10.0.100.0       n/a       FULL               10.11.2.1          vlan1102

Core-2#
```

<———————————— Begins ISP1 configuration ————————————>

Server Switch (via PC-1)

8. Using Putty, open an SSH session to the **Server-Switch.**

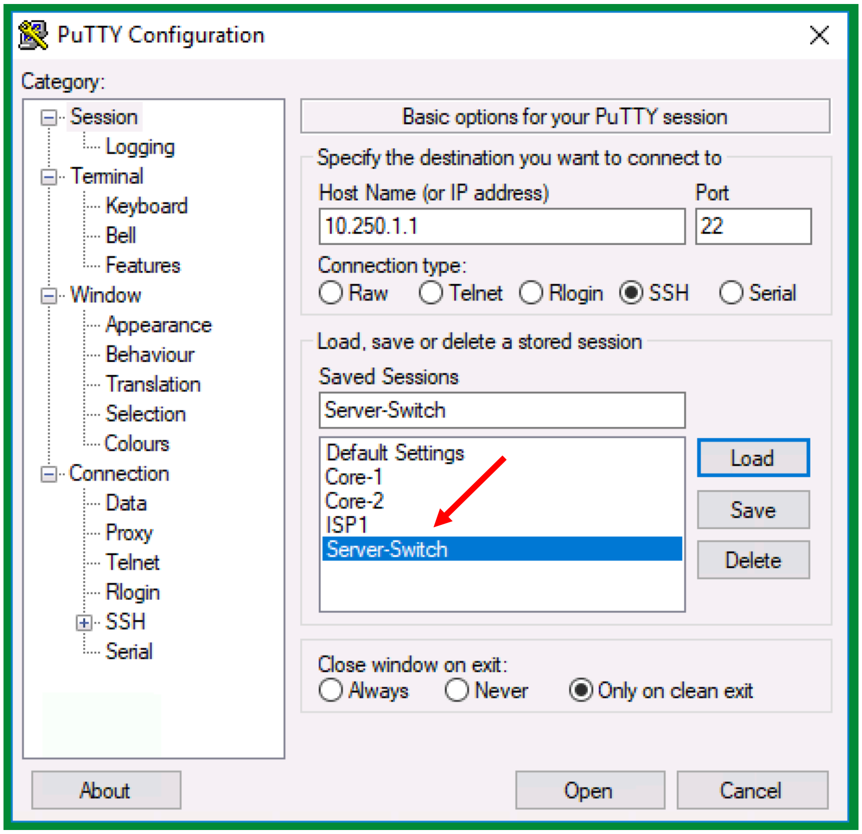

Host Name (or IP address): 10.250.1.1

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Saved Sessions: Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 11.2-21: SSH to Server Switch

9. Login using **cxfX/aruba123.**

NOTE: Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

10. Change interface **VLAN X01** and **X02**'s network type to **point-to-point**.

```
P52-6300-AB(config)# interface vlan X01
P52-6300-AB(config-if-vlan)# ip ospf network point-to-point
P52-6300-AB(config-if-vlan)# interface vlan X02
P52-6300-AB(config-if-vlan)# ip ospf network point-to-point
P52-6300-AB(config-if-vlan)# end
```

11. Close the putty session.

<——————————— Ends Server Switch configuration ———————————>

Core-2 (via PC-1)

12. Move back to Core-2.
13. Inspect **area X**'s Link State Database.

```
Core-2# show ip ospf lsdb vrf TABLE-X area X
OSPF Router with ID (10.11.100.2) (Process ID 11 VRF TABLE-11)
==============================================================

Router Link State Advertisements (Area 0.0.0.11)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
---------------------------------------------------------------------------
10.0.100.0      10.0.100.0      14        0x8000004f 0x000063c9     4
10.11.100.1     10.11.100.1     14        0x80000090 0x0000c386     6
10.11.100.2     10.11.100.2     502       0x80000094 0x0000d76e     6

Core-2#
```

How many LSAs do you have?

____

What type are they?

____

14. Compare the output with what you saw in Task 2 Step 15.

How many LSAs were suppressed and what kind were they?

____

15. Inspect the routing table and focus on OSPF prefixes.

```
Core-2# show ip route ospf vrf TABLE-X
```

```

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.0.100.0/32, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf
10.11.1.0/30, vrf TABLE-11
        via  10.11.2.1,  [110/200],  ospf
        via  10.11.0.1,  [110/200],  ospf
10.254.1.0/24, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf

Core-2#
```

Is there any prefix missing?

Task 4: Make Router ID Routable.

Objectives

Due to the uniqueness of the OSPF router ID within an Autonomous system, It is sometimes useful to use it as a system IP address in order to quickly check availability of the system by pinging it, or in the case of devices that do not support a management interface, point to that IP address whenever management is required.

Steps

PC-3

1. Access PC-3.
2. Ping the Router ID of Core-1 and Core-2.

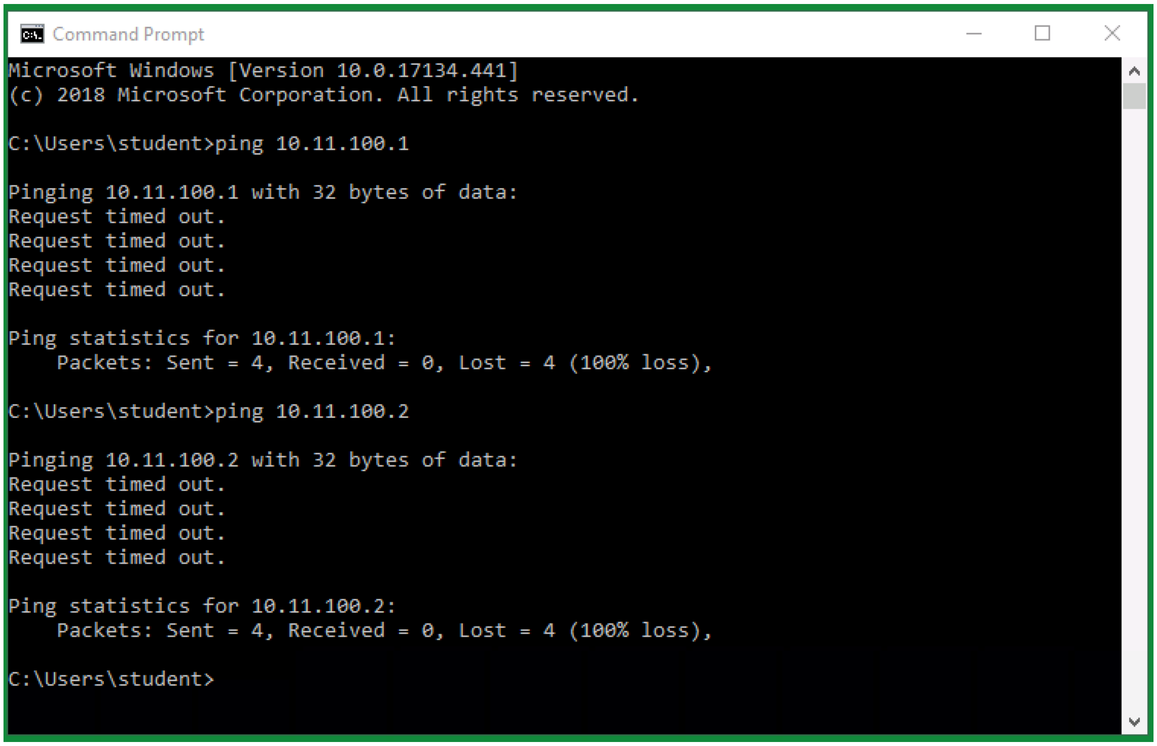

```
Microsoft Windows [Version 10.0.17134.441]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\Users\student>ping 10.11.100.1

Pinging 10.11.100.1 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 10.11.100.1:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>ping 10.11.100.2

Pinging 10.11.100.2 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 10.11.100.2:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\Users\student>
```

Figure 11.2-22: Pings timeouts.

Were the pings successful?

Why?

ANSWER: You did not get a response because the OSPF Router ID is not a system IP address. However, you can make it one by assigning it to a loopback interface.

Core-1 (via PC-1)

3. Move to Core-1
4. Create interface **loopback X** and map it to vrf **TABLE-X**.

```
Core-1# configure terminal
Core-1(config)# interface loopback X
Core-1(config-loopback-if)# vrf attach TABLE-X
```

5. Use the Router ID (**10.X.100.1**) as the loopback IP address, then enable OSPF **process X area X** in that logical interface.

```
Core-1(config-loopback-if)# ip address 10.X.100.1/32
Core-1(config-loopback-if)# ip ospf X area X
Core-1(config-loopback-if)# end
```

6. Use the **show ip ospf lsdb** command for validating the new number of links announced on Core-1's LSA.

```
Core-1# show ip ospf lsdb vrf TABLE-X router lsid 10.X.100.1
OSPF Router with ID (10.11.100.1) (Process ID 11 VRF TABLE-11)
===============================================================

Router Link State Advertisements (Area 0.0.0.11)
-------------------------------------------------
```

```
LSID            ADV Router      Age       Seq#        Checksum       Link Count
-------------------------------------------------------------------------------
10.11.100.1     10.11.100.1     240       0x80000073  0x00008836     7

Core-1#
```

NOTE: Last time (in task 3 step 13), Core-1's Router LSA had 6, after adding interface loopback X, the LSA has increased to 7.

PC-3

7. Move back to PC-3.
8. Ping the Router ID of Core-1 now. Ping should be successful.

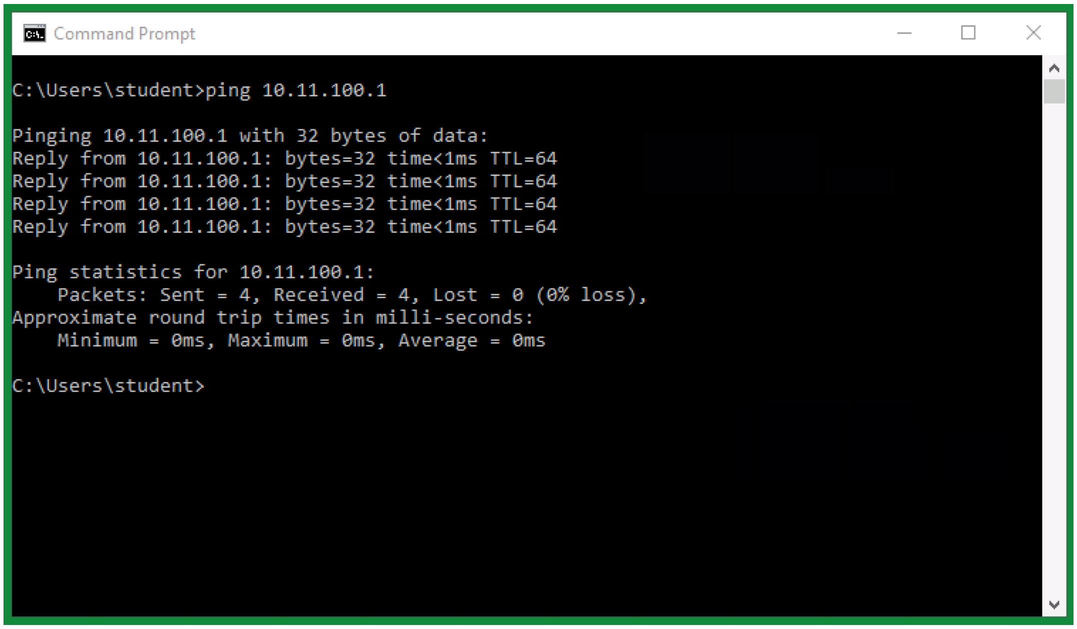

```
C:\Users\student>ping 10.11.100.1

Pinging 10.11.100.1 with 32 bytes of data:
Reply from 10.11.100.1: bytes=32 time<1ms TTL=64
Reply from 10.11.100.1: bytes=32 time<1ms TTL=64
Reply from 10.11.100.1: bytes=32 time<1ms TTL=64
Reply from 10.11.100.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.11.100.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\student>
```

Figure 11.2-23: Ping successful.

Core-2 (via PC-1)

9. Move to Core-2
10. Repeat steps 4 and 5.

```
Core-2# configure terminal
Core-2(config)# interface loopback X
Core-2(config-loopback-if)# vrf attach TABLE-X
```

```
Core-2(config-loopback-if)# ip address 10.X.100.2/32
Core-2(config-loopback-if)# ip ospf X area X
Core-2(config-loopback-if)# end
```

11. Display the OSPF routing table. You should see Core-1's Router ID value listed in the output.

```
Core-2# show ip route ospf vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.0.100.0/32, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf
10.11.1.0/30, vrf TABLE-11
        via  10.11.2.1,  [110/200],  ospf
        via  10.11.0.1,  [110/200],  ospf
10.11.100.1/32, vrf TABLE-11
        via  10.11.0.1,  [110/100],  ospf
10.254.1.0/24, vrf TABLE-11
        via  10.11.2.1,  [110/125],  ospf

Core-2#
```

Task 5: Default Network Injection.

Objectives

In lab 9, you configured static floating routes to enable redundancy while running load sharing across both internet links. You warned your customer that this solution may lead to potential Layer 3 loops if both ISPs go down.

Now that OSPF is in place, the injection of a default route through the protocol is possible in both Core switches. That will replace the floating one. Since OSPF has a local Administrative Distance of 110 and static routing has 1, this newly injected prefix will remain ignored unless the main 0.0.0.0/0 static entry vanishes after a link failure.

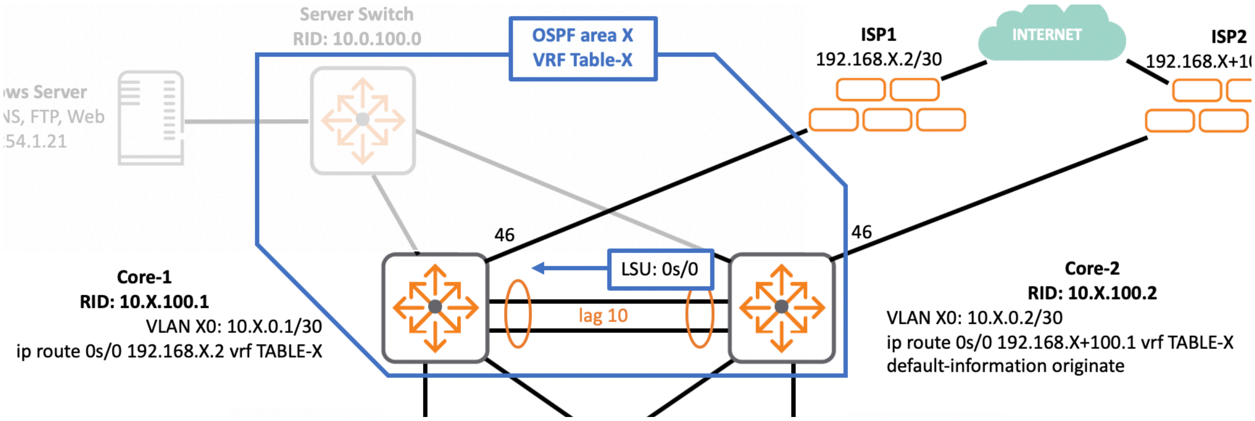

- **Figure 11.2-24: injection default route**

The main advantage of this method versus floating routes, is that Core-2 will not send this particular Link State Update if the default prefix is not present in the VRF routing table. This means if Core-2 loses its main internet link and the static route goes down, the OSPF prefix will be withdrawn.

This mechanism makes a Layer 3 loop impossible even if enabled on both Core switches.

NOTE: Default route injection uses External LSAs (LSA type 5). These LSAs are covered in more detail in the Implementing AOS-CX Switches training course.

In this task you will first remove floating routes and replace them with OSPF default route injection, then you see what happens.

Steps

Core-2 (via PC-1)

1. Open a SSH session to Core-2.
2. Move to **OSPF process X** for **TABLE-X** vrf and inject the default route.

```
Core-2# configure terminal
Core-2(config)# router ospf X vrf TABLE-X
Core-2(config-ospf-11)# default-information originate
Core-2(config-ospf-11)# exit
```

Core-1 (via PC-1)

3. Open a SSH session to Core-1.
4. Look for default routes configuration lines mapped to your VRF.

```
Core-1# show running-config | begin 0 TABLE-X | include "ip route"
ip route 0.0.0.0/0 192.168.11.2 vrf TABLE-11
ip route 0.0.0.0/0 10.11.0.2 distance 10 vrf TABLE-11
Core-1#
```

TIP: When running show command filtering tools, the matching string is typically a single word, however you can match multiple words if you quote them all between "" characters as in example above, where we are looking for lines that contain the "ip route" string.

5. Remove the default floating route.

```
Core-1# configure terminal
Core-1(config)# no ip route 0.0.0.0/0 10.X.0.2 distance 10 vrf TABLE-X
Core-1(config)#
```

6. Display the OSPF routing table in vrf **TABLE-X**.

```
Core-1(config)# show ip ospf route vrf TABLE-X
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 11 VRF TABLE-11, Routing Table
-----------------------------------------------

Total Number of Routes : 6

0.0.0.0/0          (E2)
     via 10.11.0.2 interface vlan110, cost 1 distance 110
10.0.100.0/32      (E1)
     via 10.11.1.2 interface vlan1101, cost 125 distance 110
10.11.2.0/30       (i) area: 0.0.0.11
     via 10.11.1.2 interface vlan1101, cost 200 distance 110
10.11.12.0/24      (i) area: 0.0.0.11
     via 10.11.1.2 interface vlan1101, cost 300 distance 110
10.11.100.2/32     (i) area: 0.0.0.11
     via 10.11.1.2 interface vlan1101, cost 200 distance 110
10.254.1.0/24      (E1)
     via 10.11.1.2 interface vlan1101, cost 125 distance 110

Core-1(config)#
```

Is there any default route learned by the protocol?

____

7. Look for the **0.0.0.0/0** prefix in vrf **TABLE-X** routing table.

```
Core-1(config)# show ip route 0.0.0.0 vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  192.168.11.2,  [1/0],  static

Core-1(config)#
```

How was the prefix learned?

____

Who is the next hop?

8. Disable **VLAN X91**.

```
Core-1(config)# interface vlan X91
Core-1(config-if-vlan)# shutdown
```

9. Repeat step 7.

```
Core-1(config-if-vlan)# show ip route 0.0.0.0 vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  10.11.0.2,  [110/1],  ospf

Core-1(config-if-vlan)#
```

How was the prefix learned?

What is the next hop?

NOTE: This new route can be used to forward traffic in case ISP1 fails.

Next you will simulate a failure on the link to ISP2 and see what happens to the injected route.

Core-2 (via PC-1)

10. Move back to Core-2.
11. Remove the default floating route.

```
Core-2(config)# no ip route 0.0.0.0/0 10.X.0.1 distance 10 vrf TABLE-X
```

12. Disable **VLAN X92**.

```
Core-2(config)# interface vlan X92
Core-2(config-if-vlan)# shutdown
Core-2(config-if-vlan)#
```

13. Confirm Core-2 has no default prefixes in the VRF Routing table.

```
Core-2(config-if-vlan)# show ip route 0.0.0.0 vrf TABLE-X

No ipv4 routes configured


Core-2(config)#
```

Now validate that default route injection stops taking place because Core-2 does not have a route injection entry in the VRF table.

Core-1 (via PC-1)

14. Move back to Core-1.
15. Look for the **0.0.0.0/0** prefix in the vrf **TABLE-X** routing table.

```
Core-1(config-if-vlan)# show ip route 0.0.0.0 vrf TABLE-X
No ipv4 routes configured


Core-1(config-if-vlan)#
```

Is there any default route in the VRF routing table?

16. Take a look into the OSPF process' routing table for vrf **TABLE-X**.

```
Core-1(config-if-vlan)# show ip ospf route 0.0.0.0/0 vrf TABLE-X
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 11 VRF TABLE-11, Routing Table for prefixes 0.0.0.0/0
---------------------------------------------------------------------

Total Number of Routes : 0

Core-1(config-if-vlan)#
```

Is there any 0s prefix?

Why?

Now restore the ISP1 link and enable the route injection in Core-1 as well as confirming that Core-2 is now learning the route via Core-1.

17. Enable interface VLAN X91 (you disabled it in step 8).

```
Core-1(config-if-vlan)# no shutdown
Core-1(config-if-vlan)# exit
```

18. Move to OSPF process **X** for vrf **TABLE-X** and inject the default route.

```
Core-1(config)# router ospf X vrf TABLE-X
Core-1(config-ospf-11)# default-information originate
Core-1(config-ospf-11)# end
```

Core-2 (via PC-1)

19. Move back to Core-2.
20. Look for the **0.0.0.0/0** prefix in the vrf **TABLE-X** routing table.

```
Core-2(config-if-vlan)# show ip route 0.0.0.0 vrf TABLE-X

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf TABLE-11
        via  10.11.0.1,  [110/1],  ospf

Core-2(config-if-vlan)#
```

Is there any default route in the VRF routing table?

What is the next hop?

21. Enable interface VLAN X92.

```
Core-2(config-if-vlan)# no shutdown
Core-2(config-if-vlan)# end
Core-2#
```

Task 6: Save Your Configurations

Objectives

Save your configuration.

Steps

Core-1 and Core-2 (via PC-1)

1. Save the current Core-1 and Core-2 switch configuration in the startup checkpoint.

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

You have completed Lab 11.2!

AOS-CX Switching Fundamentals

Lab 12.1: Create a Virtual Switching Framework Stack

Overview

it has been one year since BigStartup started business and increased profits are making it possible to open additional offices. This new project for additional offices begins next month and they want you to take care of the entire network deployment. This project will take several months and you might not be able to assist with Level 1 support. You suggest handing over control of the access switches to an internal staff member. He is not very experienced in networking and does not feel confident managing multiple independent switches.

In order to simplify the deployment, you plan to create a single stack of switches using a technology called Virtual Switching Framework (VSF) so he only will need to deal with one logical unit.

Objectives

After completing this lab, you will be able to:

- Create a VSF stack
- Define stack roles
- Verifying VSF topology
- Configure distributed Link Aggregation

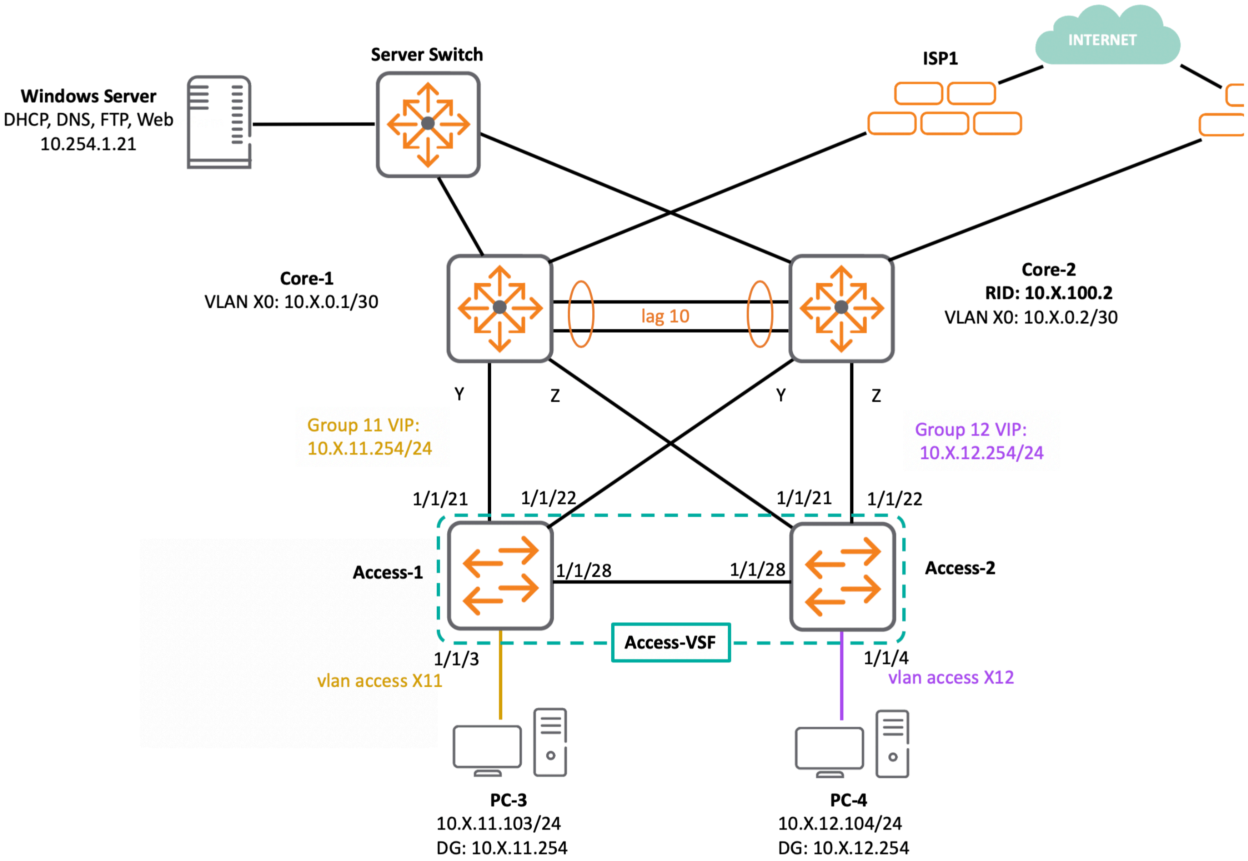

Figure 12.1-1: Lab Topology

Task 1: Deploy a VSF Stack.

Objectives

You are about to create a VSF stack. This involves rebooting one of the units which might affect users connected to it. Although you know the process will take no more than 5 minutes, you have requested a 30 minutes maintenance window. To further minimize the inconvenience, you have scheduled the maintenance window during lunch.

In this task, you will create a VSF stack with both Access switches using port 1/1/28. Then you will explore the stack properties and normalize the port configuration on member 2.

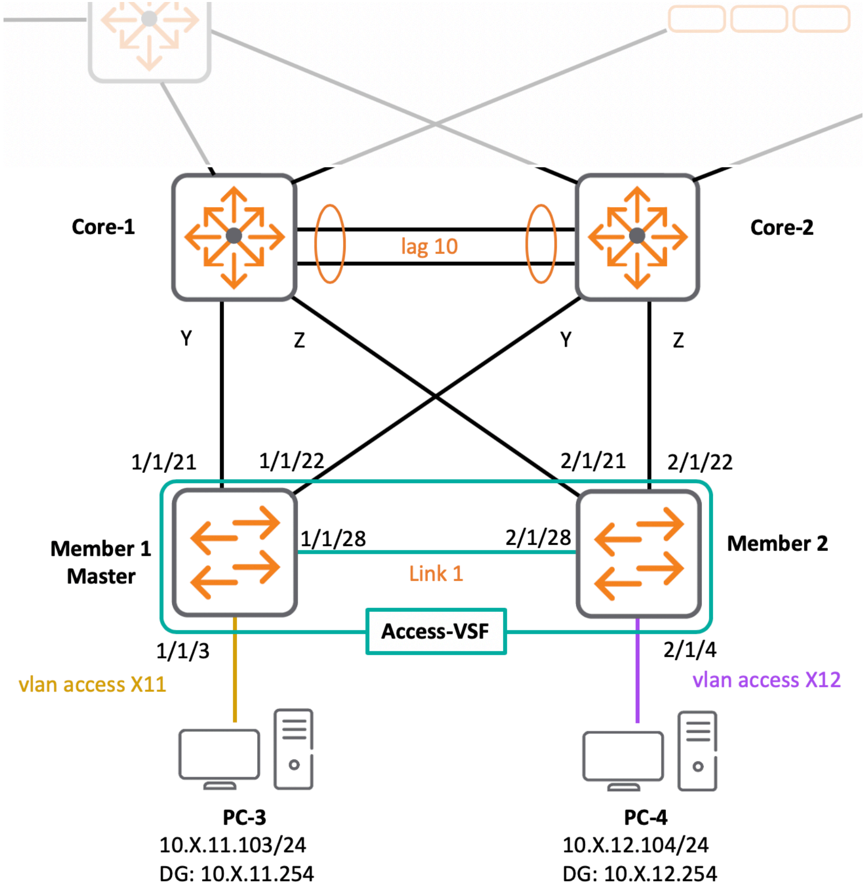

Figure 12.1-2: Access VSF

Steps

PC-4

1. Open a console session to PC-4.
2. Run a continuous ping to **8.8.8.8**. Ping should be successful.

Access-1

3. Open a console session to Access-1.
4. Create VSF link 1 using port **1/1/28.**

```
T11-Access-1# configure terminal
T11-Access-1(config)# vsf member 1
T11-Access-1(vsf-member-1)# link 1 1/1/28
T11-Access-1(vsf-member-1)# exit
```

Access-2

5. Open a console session to Access-2.
6. Create VSF link 1 using port **1/1/28.**

```
T11-Access-2(config)# vsf member 1
T11-Access-2(vsf-member-1)# link 1 1/1/28
T11-Access-2(vsf-member-1)# exit
```

7. Renumber the switch to vsf member 2. You will be prompted to save configuration and reboot the unit. Answer “**y**”.

```
T11-Access-2(config)# vsf renumber-to 2
This will save the VSF configuration and reboot the switch.
Do you want to continue (y/n)? y
```

The system will reboot and be back online after a few minutes.

8. Login with **admin** and no password (leave empty).

```
T11-Access-2 login: admin
Password:

member#
```

What is the new prompt shown in the switch's CLI?

Access-1

9. Move back to Access-1.
10. Run the "**show vsf**" command.

```
T11-Access-1(config)# show vsf

MAC Address                : 88:3a:30:98:30:00
Secondary                  :
Topology                   : Chain
Status                     : No Split
Split Detection Method     : None


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   88:3a:30:98:30:00   JL668A         Master
2   88:3a:30:97:a4:40   JL668A         Member
T11-Access-1(config)#
```

What is the Stack's MAC address?

What is the topology used in the stack?

How many members are part of the stack?

Does the stack MAC address matches any of the member's?

Whose?

What is status (role) of Member 1?

What is status (role) of Member 2?

11. Run the detailed version of the output.

```
T11-Access-1(config)# show vsf detail
VSF Stack
        MAC Address               : 88:3a:30:98:30:00
        Secondary                 :
        Topology                  : chain
        Status                    : No Split
        Split Detection Method    : None
        Software Version          : FL.10.04.0003

        Name                      : Aruba-VSF-6300
        Contact                   :
        Location                  :

Member ID                         : 1
```

```
        MAC Address                 : 88:3a:30:98:30:00
        Type                        : JL668A
        Model                       : 6300F 24-port 1GbE and 4-port SFP56 Switch
        Status                      : Master
        ROM Version                 : FL.01.05.0003
        Serial Number               : SG90KN70HX
        Uptime                      : 4 days, 10 hours, 4 minutes
        CPU Utilization             : 1%
        Memory Utilization          : 17%
        VSF Link 1                  : Up, connected to peer member 2, link 1
        VSF Link 2                  :

Member ID                           : 2
        MAC Address                 : 88:3a:30:97:a4:40
        Type                        : JL668A
        Model                       : 6300F 24-port 1GbE and 4-port SFP56 Switch
        Status                      : Member
        ROM Version                 : FL.01.05.0003
        Serial Number               : SG90KN70HB
        Uptime                      : 2 minutes
        CPU Utilization             : 1%
        Memory Utilization          : 8%
        VSF Link 1                  : Up, connected to peer member 1, link 1
        VSF Link 2                  :

T11-Access-1(config)#
```

What is the switch type (part number) of both members?

____

What is the switch type (Model) of both members?

____

____

What is the CPU and memory utilization of Member 1?

____

What is the CPU and memory utilization of Member 2?

____

12. Use the **show vsf topology** for looking at logical connections between members.

```
T11-Access-1(config)# show vsf topology
          Mstr
 +---+    +---+
 | 2 |1==1| 1 |
 +---+    +---+

T11-Access-1(config)#
```

What is the logical link that connect both units?

____

13. Run the **show vsf link** for displaying the physical port members of logical link 1.

```
T11-Access-1(config)# show vsf link

 VSF Member 1

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ ---------------------------
1    up         2       1      1/1/28


 VSF Member 2

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ ---------------------------
1    up         1       1      2/1/28

T11-Access-1(config)#
```

What ports are used in Member 1 for creating VSF link 1?

____

What ports are used in Member 2 for creating VSF link 1?

____

Both members are now part of the same logical stack. They share the same control plane and management plane, although data plane is distributed among them. It means that the physical interfaces of both units can be managed by the Master.

14. Run the **show interface brief** and confirm you can see ports of both members.

```
T11-Access-1(config)# show interface brief
--------------------------------------------------------------------------------
----
Port      Native  Mode   Type           Enabled Status  Reason
Speed
          VLAN
(Mb/s)
--------------------------------------------------------------------------------
1/1/1     1112    access 1GbT           yes     up                        1000
1/1/2     1       access 1GbT           yes     down    Administratively down  --
1/1/3     1111    access 1GbT           yes     up                        1000
1/1/4     1       access 1GbT           no      down    Administratively down  --
                                            ←---- output omitted ---→
1/1/21    1       trunk  1GbT           yes     up                        1000
1/1/22    1       trunk  1GbT           yes     up                        1000
1/1/23    1       access 1GbT           no      down    Administratively down  --
1/1/24    1       access 1GbT           no      down    Administratively down  --
1/1/25    1       access SFP+DAC1       no      down    Administratively down  --
1/1/26    1       access SFP+DAC1       no      down    Administratively down  --
1/1/27    1       access SFP+DAC1       no      down    Administratively down  --
1/1/28    --      routed SFP+DAC1       yes     up                       10000
2/1/1     1       access 1GbT           yes     down    Waiting for link       --
2/1/2     1       access 1GbT           yes     down    Waiting for link       --
2/1/3     1       access 1GbT           yes     down    Waiting for link       --
2/1/4     1       access 1GbT           yes     up                        1000
                                            ←---- output omitted ---→
2/1/21    1       access 1GbT           yes     up                        1000
2/1/22    1       access 1GbT           yes     up                        1000
2/1/23    1       access 1GbT           yes     down    Waiting for link       --
2/1/24    1       access 1GbT           yes     down    Waiting for link       --
2/1/25    1       access SFP+DAC1       yes     up                       10000
2/1/26    1       access SFP+DAC1       yes     down    Waiting for link       --
2/1/27    1       access SFP+DAC1       yes     down    Waiting for link       --
2/1/28    --      routed SFP+DAC1       yes     up                       10000
vlan1     --             --             yes     up                           --
T11-Access-1(config)#
```

Can you see ports of member 1 and member 2?

What is the mode of interfaces used for the VSF link?

> **ANSWER:** These interfaces lost their previous configuration, moved to routed ports and are now exclusively used for VSF. Due their routed mode properties layer 2 loops cannot be created through them.

What VLANs are assigned to ports 1/1/1 and port 1/1/3 (PC-1 and PC-3)?

What VLAN is assigned to port 2/1/4 (PC-4)?

What is the port mode of interfaces 1/1/21 and 1/1/22 (uplinks of Member1)?

What is the port mode of interfaces 2/1/21 and 2/1/22 (uplinks of Member2)?

PC-4

15. Move back to PC-4.

Is the ping still going?

____

> **NOTICE:** When member 2 came back from rebooting and joined the stack, it lost its previous configuration, wiping the ports' settings out and putting them in default values. This process is obviously affecting PC-4 who can no longer access internet.
>
> You realize you only have 10 minutes left before the maintenance window is over. So, you better hurry up and restore the configuration on those ports!
>
> Do not panic! You do not have to create the VLANs or Spanning Tree configuration all over again, they are already part of the global VSF stack configuration that Member 1 manages. The only thing you must do is to provision the ports properly.

Access-1

16. Move back to Access-1.
17. Disable all Member 2's ports but the VSF connections?

```
T11-Access-1(config)# interface 2/1/1-2/1/27
T11-Access-1(config-if-<2/1/1-2/1/27>)# shutdown
T11-Access-1(config-if-<2/1/1-2/1/27>)# exit
```

18. Enable Member 2's uplinks to Core-1 and Core-2 and allow VLANs X11 and X12 across interfaces 2/1/21 and 2/1/22.

```
T11-Access-1(config)# interface 2/1/21-2/1/22
T11-Access-1(config-if-<2/1/21-2/1/22>)# no shutdown
T11-Access-1(config-if-<2/1/21-2/1/22>)# vlan trunk allowed X11-X12
T11-Access-1(config-if-<2/1/21-2/1/22>)# exit
```

19. Enable the port that connects to **PC-4** (**2/1/4**), then make it member of **VLAN X12**.

```
T11-Access-1(config)# interface 2/1/4
T11-Access-1(config-if)# no shutdown
T11-Access-1(config-if)# vlan access X12
T11-Access-1(config-if)# exit
```

Well done! You have restored connectivity in record time! Now that the urgency is over, you can change the hostname of the system to something more appropriate.

20. Change the hostname to **TX-Access-VSF**.

```
T11-Access-1(config)# hostname TX-Access-VSF
```

PC-4

21. Move back to PC-4.

Is the ping working now?

____

22. Stop the ping.

Task 2: Configure Distributed Link Aggregation.

Objectives

Right now, the stack is up and running. However, because of your Spanning Tree knowledge, you know that only two out of the four uplinks are actively in use: 1/1/21 is root port for Instance 1 and alternate for Instance 2 while 1/1/22 is root port for instance 2 and alternate on Instance 1. The other two uplinks 2/1/21 and 2/1/22 are alternate of both instances.

Therefor you must complete the deployment by configuring Link aggregation between the Stack and both Cores.

You will first create lag X1 in both the VSF stack and Core-1. Then you will create lag X2 in Core-2 and the VSF stack.

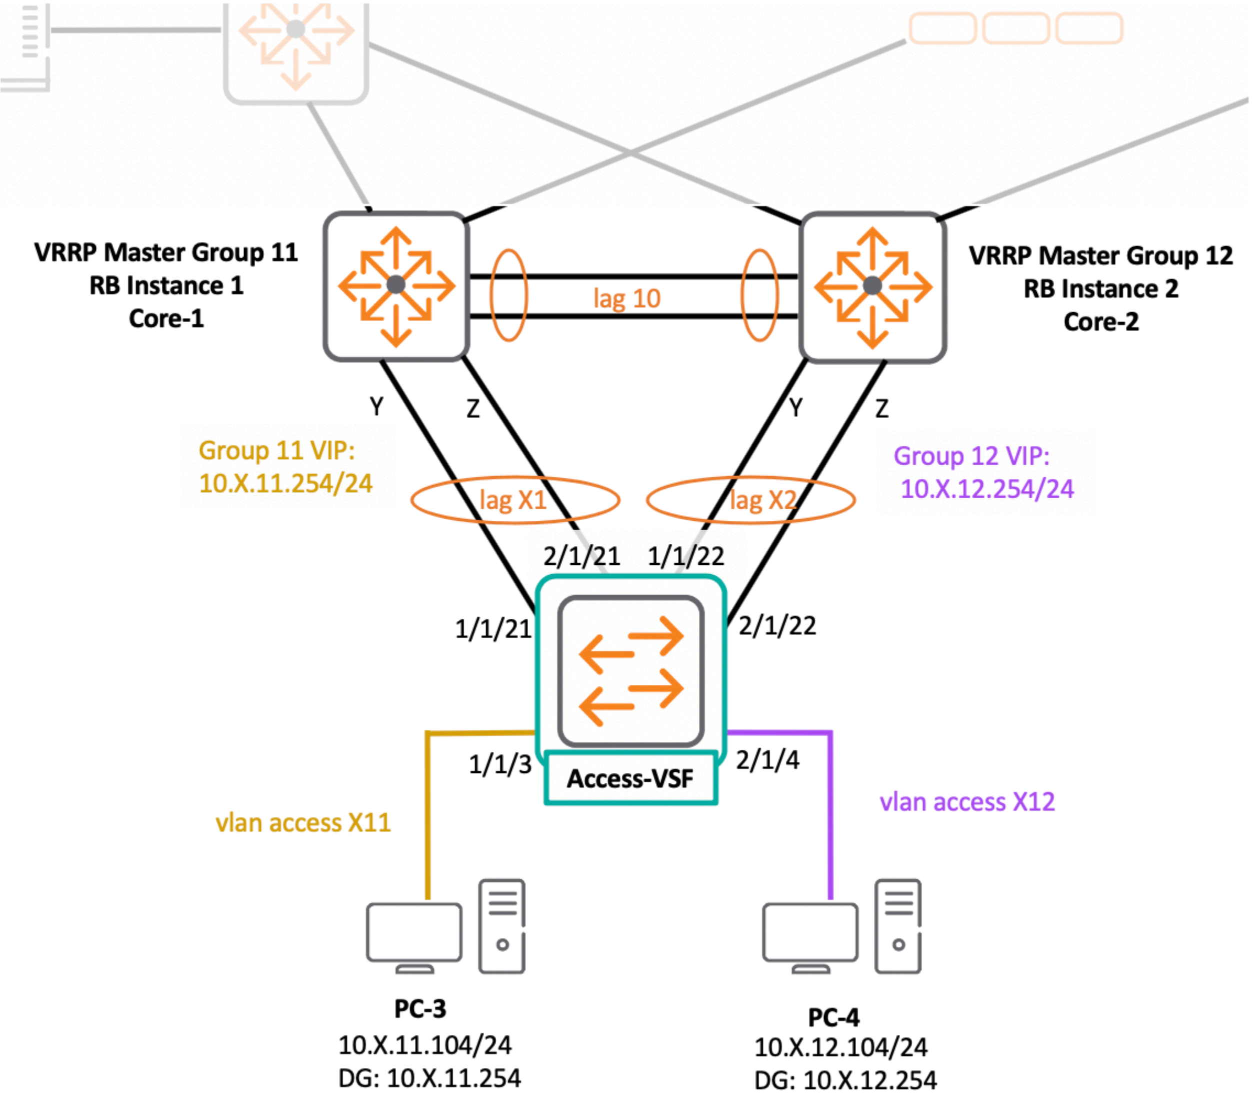

Figure 12.1-3: LACP Topology

Steps

PC-3

1. Access PC-3.
2. Run a continuous ping to PC-4 **(10.X.12.104)**. Ping should be successful.

Access-VSF: Member 2

3. Open a console session to Access-VSF: Member 2 (formerly known as Access-2).
4. Hit the "**?**" question mark. You will get the help as the output.

```
member# ?
  diagnostics        Change diagnostic commands availability
  exit               Exit current mode and change to previous mode
  list               Print command list
  member             VSF member selection
  no                 Negate a command or set its defaults
  page               Enable page break
  show               Show running system information
  start-shell        Start Bash shell
  vsf-factory-reset  Clear all VSF configurations and boot as the primary
                     switch
member#
```

5. Type "**show**" followed by "**?**" question mark. You will get the "show" command's help as the output.

```
member# show ?
  session-timeout  Idle session timeout in minutes
  tech             Display output of a predefined command sequence used by
                   technical support
  version          Displays switch version
  vsf              Show VSF information
member#
```

Are the available commands and options the same that you would see in the Master or a non-stacked switch?

____

____

6. Run the **member 1** command, this will take you to Member 1's (the master) CLI.

```
member# member 1
T11-Access-VSF#
```

7. Create **lag X1** with the following settings:
    a) Description: **TO_CORE-1**

b) Allowed VLANs: **X11 and X12**.

c) LACP rate: **fast**

d) LACP mode: **active**

e) Enabled: **yes**

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface lag X1
T11-Access-VSF(config-lag-if)# description TO_CORE-1
T11-Access-VSF(config-lag-if)# vlan trunk allowed X11-X12
T11-Access-VSF(config-lag-if)# lacp mode active
T11-Access-VSF(config-lag-if)# lacp rate fast
T11-Access-VSF(config-lag-if)# no shutdown
T11-Access-VSF(config-lag-if)# exit
```

8. Associate ports **1/1/21** and **2/1/21** to **lag X1**.

```
T11-Access-VSF(config)# interface 1/1/21
T11-Access-VSF(config-if)# lag X1
T11-Access-VSF(config-if)# exit
T11-Access-VSF(config)# int 2/1/21
T11-Access-VSF(config-if)# lag X1
T11-Access-VSF(config-if)# exit
```

PC-3

9. Move back to PC-3.

Is the ping still running?

____

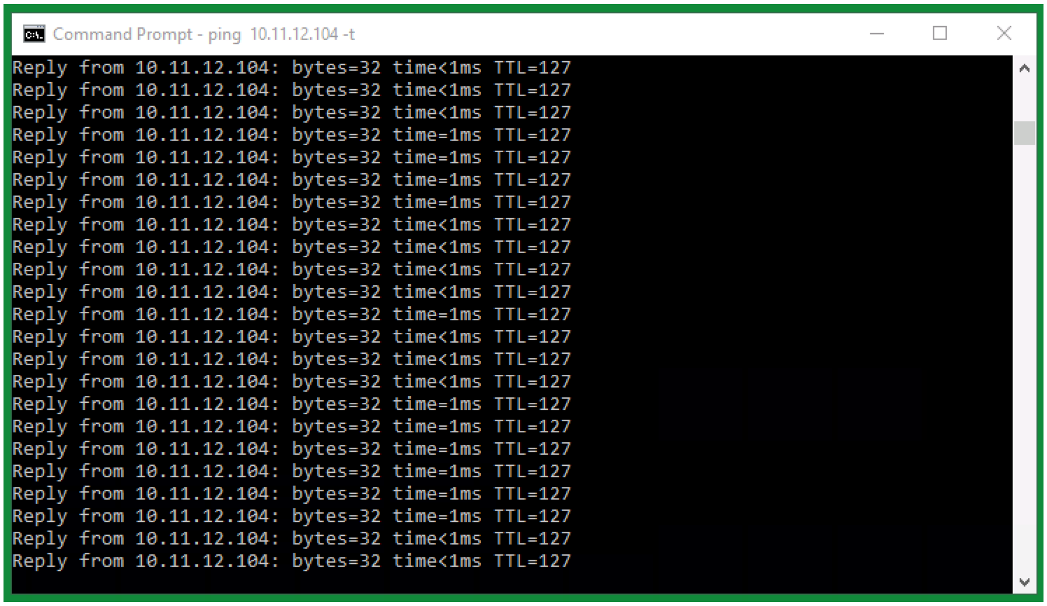

```
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
```

Figure 12.1-4: Ping to Pc4

Core-1 (via PC-1)

10. Open an SSH session to Core-1. Login using **cxfX/aruba123.**

NOTE: Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

11. Create **lag X1** with the following settings:
    a) Description: **TO_TX-ACCESS-VSF**
    b) Routing: **no**
    c) Allowed VLANs: **X11 and X12**.
    d) LACP rate: **fast**
    e) LACP mode: **active**
    f) Enabled: **yes**

```
Core-1# configure terminal
Core-1(config)# interface lag X1
Core-1(config-lag-if)# description TO_TX-ACCESS-VSF
Core-1(config-lag-if)# no routing
Core-1(config-lag-if)# vlan trunk allowed X11-X12
```

```
Core-1(config-lag-if)# lacp mode active
Core-1(config-lag-if)# lacp rate fast
Core-1(config-lag-if)# no shutdown
```

12. Associate ports **1/1/Y** and **1/1/Z** to **lag X1**.

> **TIP:** If you do not remember what ports your downlinks are, then please refer to "Lab - 4.3 - Add a Core Switch to the Topology - Task 1" for getting the right port numbers.

```
Core-1(config)# interface 1/1/Y
Core-1(config-if)# lag X1
Core-1(config-if)# exit
Core-1(config)# interface 1/1/Z
Core-1(config-if)# lag X1
Core-1(config-if)# exit
```

Core-2 (via PC-1)

13. Open an SSH session to Core-2.
14. Repeat steps 11 and 12, creating **lag X2** instead.

> **NOTE:** Replace the highlighted "X" for your student table number.

```
Core-2# configure terminal
Core-2(config)# interface lag X2
Core-2(config-lag-if)# description TO_TX-ACCESS-VSF
Core-2(config-lag-if)# no routing
Core-2(config-lag-if)# vlan trunk allowed X11-X12
Core-2(config-lag-if)# lacp mode active
Core-2(config-lag-if)# lacp rate fast
Core-2(config-lag-if)# no shutdown
Core-1(config-lag-if)# exit
```

```
Core-2(config)# interface 1/1/Y
Core-2(config-if)# lag X2
Core-2(config-if)# exit
Core-2(config)# interface 1/1/Z
Core-2(config-if)# lag X2
Core-2(config-if)# exit
```

PC-3

15. Move back to PC-3.

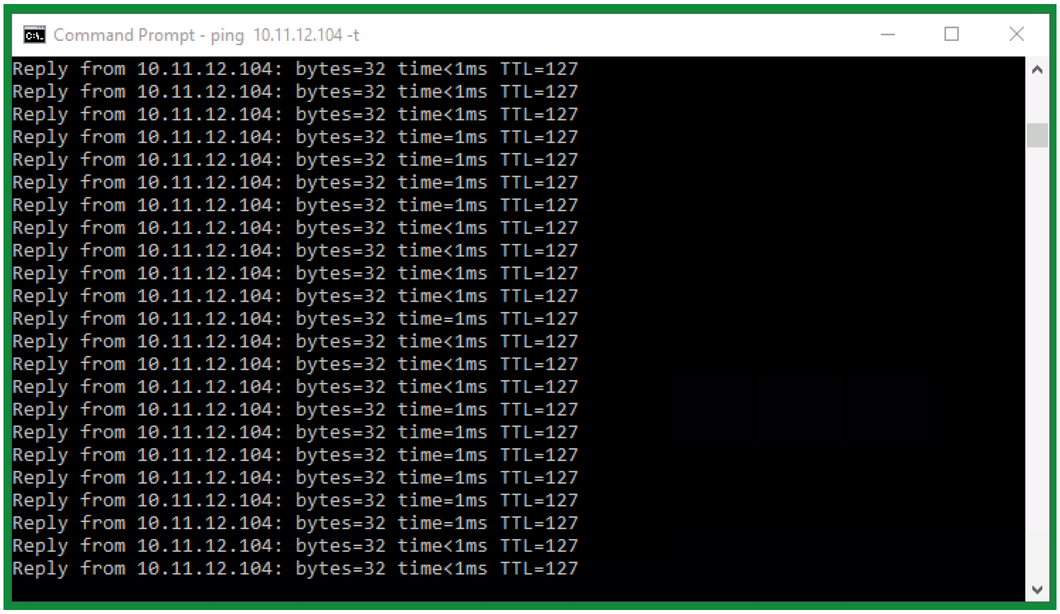

Command Prompt - ping 10.11.12.104 -t

```
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
```

Figure 12.1-5: Ping Pc4

Is the ping still running?

____

Access-VSF: Member 1

16. Move to Member 1.

17. Repeat step 14 using TO_CORE-2 as description and mapping the lag to ports **1/1/22** and **2/1/22** instead.

```
T11-Access-VSF(config)# interface lag X2
T11-Access-VSF(config-lag-if)# description TO_CORE-2
T11-Access-VSF(config-lag-if)# vlan trunk allowed X11-X12
T11-Access-VSF(config-lag-if)# lacp mode active
T11-Access-VSF(config-lag-if)# lacp rate fast
T11-Access-VSF(config-lag-if)# no shutdown
```

```
T11-Access-VSF(config-lag-if)# exit
```

```
T11-Access-VSF(config)# interface 1/1/22
T11-Access-VSF(config-if)# lag X2
T11-Access-VSF(config-if)# exit
T11-Access-VSF(config)# interface 2/1/22
T11-Access-VSF(config-if)# lag X2
T11-Access-VSF(config-if)# end
```

18. Run the **show lacp interfaces** command, then confirm all four uplinks are UP.

```
T11-Access-VSF# show lacp interfaces | begin 8 Actor
Actor details of all interfaces:
------------------------------------------------------------------------------
Intf    Aggr        Port  Port  State   System-ID         System Aggr Forwarding
        Name        Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/21  lag111      22    1     ALFNCD  88:3a:30:98:30:00 65534  111  up
2/1/21  lag111      86    1     ALFNCD  88:3a:30:98:30:00 65534  111  up
1/1/22  lag112      23    1     ASFNCD  88:3a:30:98:30:00 65534  112  up
2/1/22  lag112      87    1     ASFNCD  88:3a:30:98:30:00 65534  112  up
T11-Access-VSF#
```

19. Use the **show spanning-tree mst 1** command for validating lag X1 is root and lag X2 is alternate.

```
T11-Access-VSF# show spanning-tree mst 1 | begin 60 Port | exclude Disabled
               Port:lag111, Cost:20000, Rem Hops:19

Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
1/1/1          Designated     Forwarding   20000   128        point_to_point
1/1/3          Designated     Forwarding   20000   128        point_to_point
2/1/4          Designated     Forwarding   20000   128        point_to_point
lag111         Root           Forwarding   20000   64         point_to_point
lag112         Alternate      Blocking     20000   64         point_to_point

Topology change flag         : True
Number of topology changes   : 11
Last topology change occurred : 239 seconds ago

T11-Access-VSF#
```

20. Use the **show spanning-tree mst 2** command for validating lag X2 is root and lag X1 is Alternate.

```
T11-Access-VSF# show spanning-tree mst 2 | begin 60 Port | exclude Disabled
               Port:lag112, Cost:20000, Rem Hops:19

Port           Role           State        Cost    Priority   Type
-------------- -------------- ------------ ------- ---------- ----------
1/1/1          Designated     Forwarding   20000   128        point_to_point
1/1/3          Designated     Forwarding   20000   128        point_to_point
2/1/4          Designated     Forwarding   20000   128        point_to_point
lag111         Alternate      Blocking     20000   64         point_to_point
lag112         Root           Forwarding   20000   64         point_to_point

Topology change flag          : True
Number of topology changes    : 16
Last topology change occurred : 256 seconds ago

T11-Access-VSF#
```

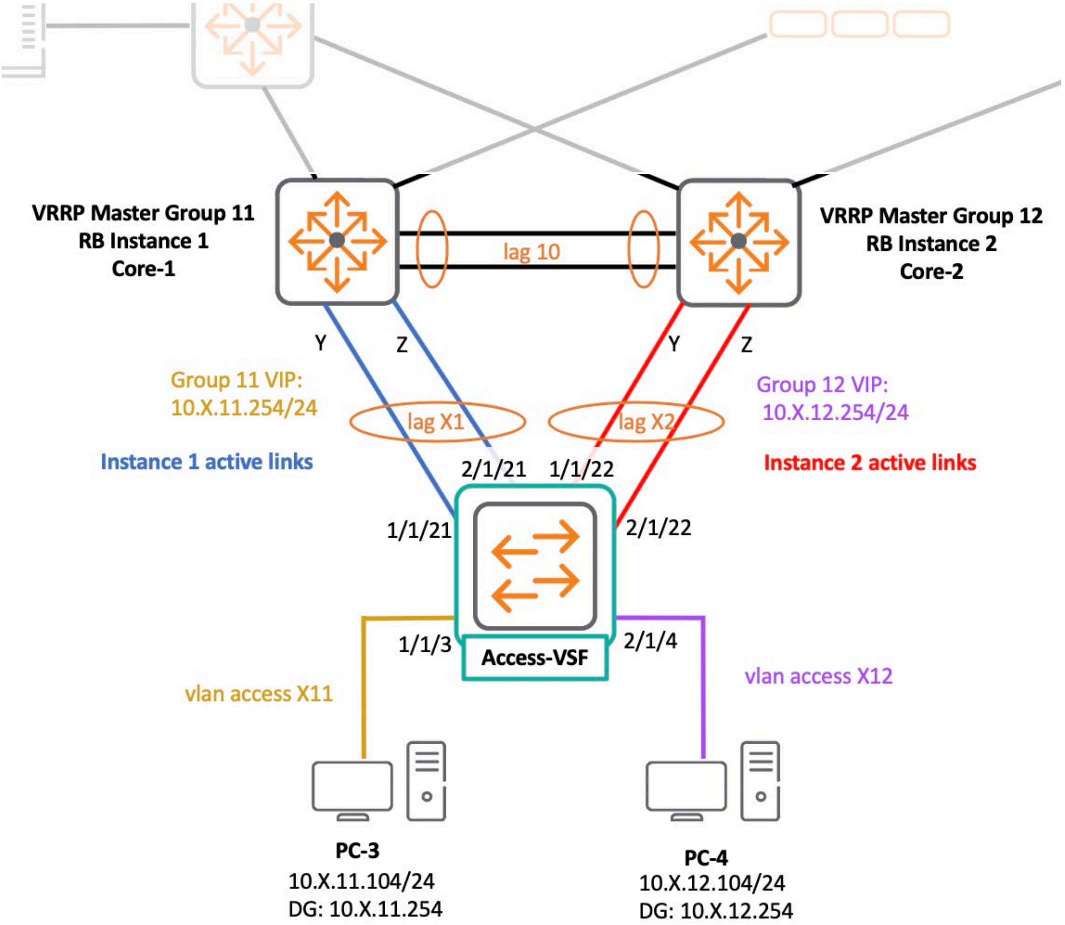

Figure 12.1-6: MST

PC-3

21. Move back to PC-3.

Is the ping still running?

____

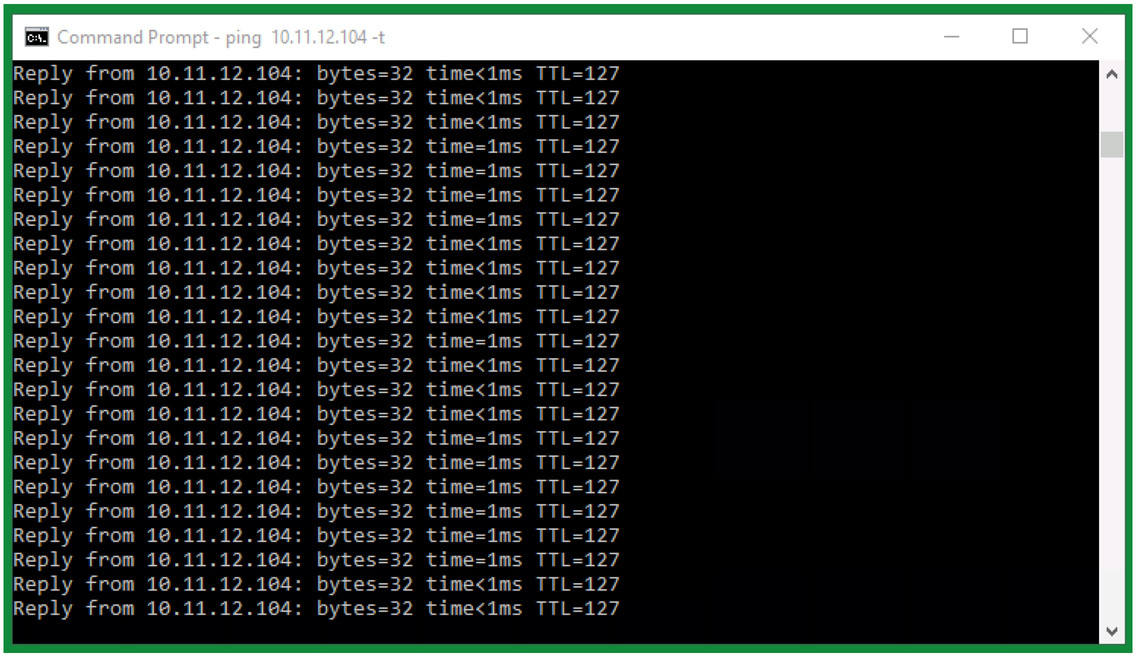

```
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time=1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
```

Figure 12.1-7: Ping to PC-4

Task 3: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

Access-VSF, Core-1 and Core-2 (via PC-1).

1. Save the current Access and Core switches' configuration in the startup checkpoint.

```
T11-Access-VSF # write memory
Configuration changes will take time to process, please be patient.
T11-Access-VSF #
Access-2#
```

```
Core-1# write memory
Configuration changes will take time to process, please be patient.
Core-1#
```

```
Core-2# write memory
Configuration changes will take time to process, please be patient.
Core-2#
```

Access-VSF

2. Backup the current Access-VSF's configuration as a custom checkpoint called **Lab12-1_final**.

```
T11-Access-VSF # copy running-config checkpoint Lab12-1_final
Configuration changes will take time to process, please be patient.
T11-Access-VSF#
```

You have completed Lab 12.1!

AOS-CX Switching Fundamentals

Lab 12.2: Maintaining the VSF Stack

Overview

After deploying VSF and centralizing both the control and management plane, the next phase is to assure there is no single point of failure that could prevent the stack from working. This is done by enabling two main features: standby member and split detection. In order to test these features BigStartup has authorized another maintenance window.

Objectives

After completing this lab, you will be able to:

- Increase the stack resiliency by adding a standby member
- Provide stack stability using split-detection
- Validate the proper performance of the features

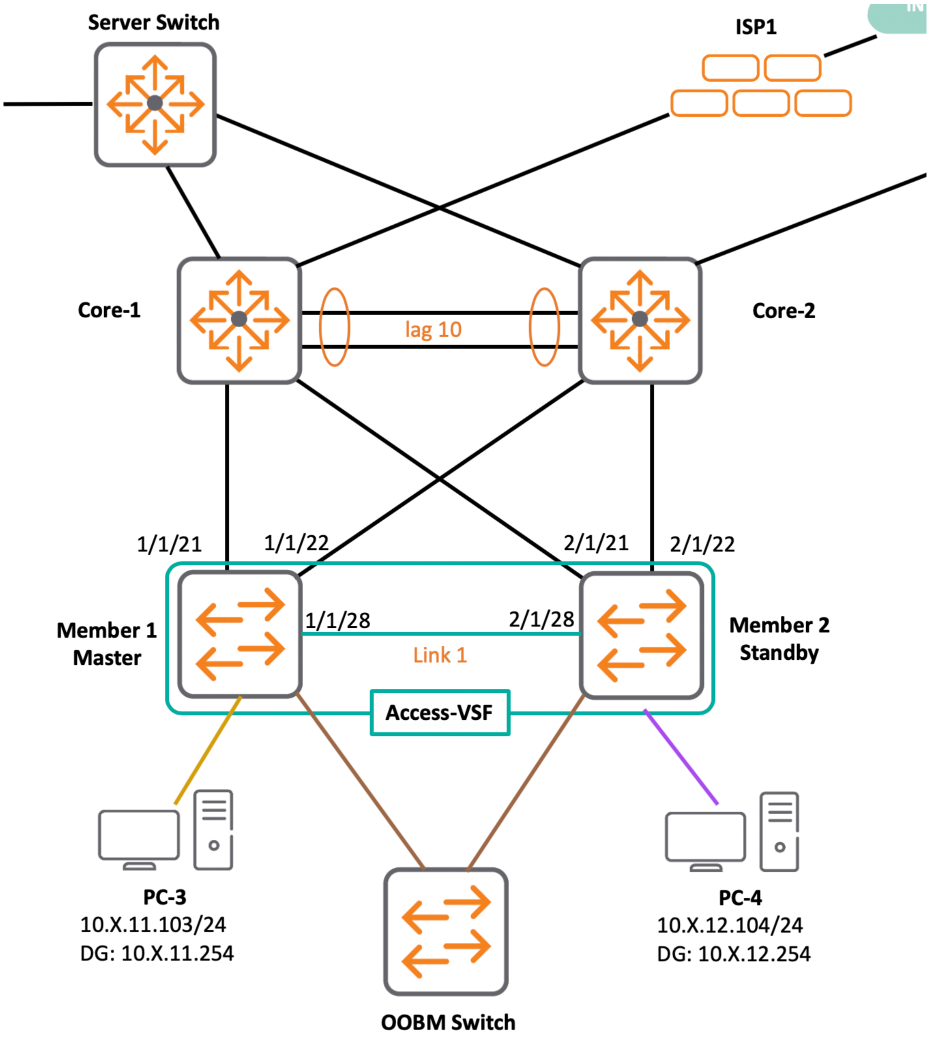

Figure 12.2-1: Lab Topology

Task 1: Secondary Member

Objectives

Once the stack is created and traffic is flowing, next step is to maintain the stack and make sure it is as stable as possible.

NOTE: Currently there is a single Master taking care of the management and control plane duties. If that switch happens to fail, then the stack will lose its main point of control and the whole stack goes down, getting stuck in the boot process as seen in console output below.

```
[  OK  ] Started PVNET namespace move script.
[FAILED] Failed to start HPE Credential Manager.
See 'systemctl status hpe-credmgr.service' for details.
[  OK  ] Reached target VSF Discovery System.
[  OK  ] Started HA Type Check Service.
[  OK  ] Reached target Check HA Target to Boot.
[  OK  ] Stopped HPE Credential Manager.
         Starting HPE Credential Manager...

[FAILED] Failed to start HPE Credential Manager.
See 'systemctl status hpe-credmgr.service' for details.
[  OK  ] Stopped HPE Credential Manager.
         Starting HPE Credential Manager...
```

In order to break this loop, the only alternative is to invoke recovery mode pressing the **[Ctrl]**+**[C]** key sequence, taking the member(s) into "**recovery**" mode.

```
**********************************************************************
WARNING! Entering emergency support login mode. This mode is for
support use only and the system will not be fully operational.
The system must be rebooted to restore full operation.
**********************************************************************
T11-Access-VSF login: admin
Password:

recovery#
```

In such case, you have to recover the master and “**reboot**” the member, otherwise you would have to set the switches into factory default using the “**vsf-factory-reset**” recovery context command and configure them all over again.

```
recovery# ?
  boot                 Reboot all or part of the system; configure default
boot
                       parameters
  copy                 Copy data or files to/from the switch
  erase                Erase device information or files
  exit                 Exit current mode and change to previous mode
  list                 Print command list
  show                 Show running system information
  start-shell          Start Bash shell
  vsf-factory-reset    Clear all VSF configurations and boot as the primary
                       switch
recovery#
```

In order to prevent this situation from happening, you can assign (in advance) the “standby” role (secondary member) to any other member of the stack. Once assigned, upon failure of the master, the standby member will take over the master role.

In this lab you will assign the standby role to Member 2 and simulate a failure on Member 1.

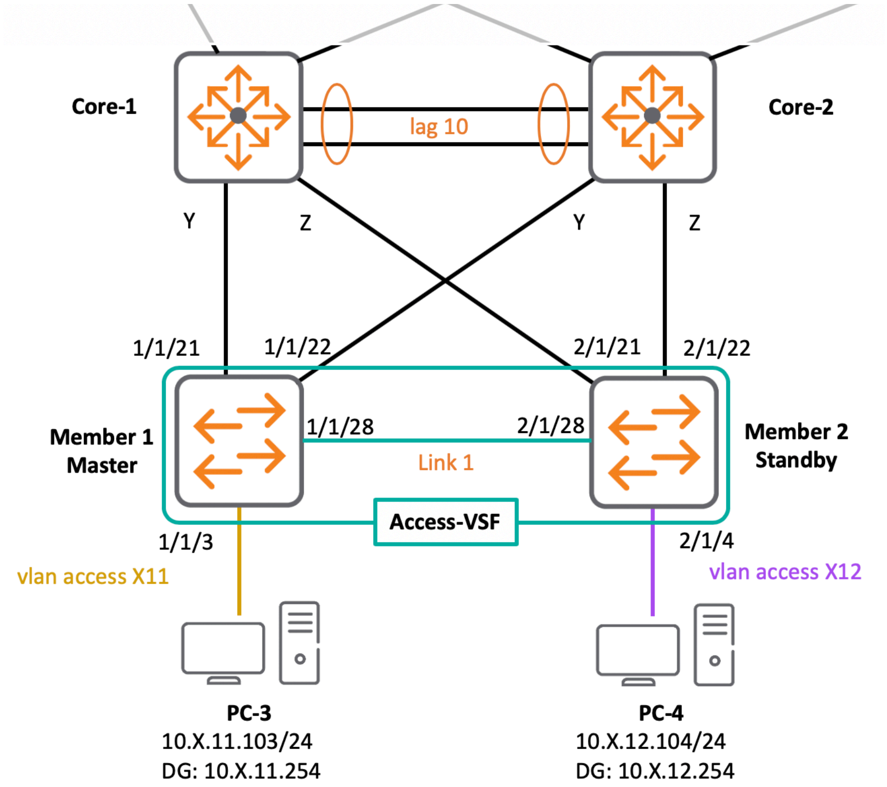

Figure 12.2-2: Lab 12.2 Task 2 Topology

Steps

Access-VSF: Member 1

1. Access Member 1’s console session.
2. Assign the standby member role. Member 2 will reboot.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# vsf secondary-member 2
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y
T11-Access-VSF(config)#
```

3. After a few minutes issue the **show vsf** and **show vsf topology** command to see the new role assigned to member 2.

```
T11-Access-VSF(config)# show vsf

MAC Address                : 88:3a:30:98:30:00
Secondary                  : 2
Topology                   : Chain
Status                     : No Split
Split Detection Method     : None


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   88:3a:30:98:30:00   JL668A         Master
2   88:3a:30:97:a4:40   JL668A         Standby
T11-Access-VSF(config)#
```

```
T11-Access-VSF(config)# show vsf topology
 Stdby    Mstr
 +---+    +---+
 | 2 |1==1| 1 |
 +---+    +---+

T11-Access-VSF(config)#
```

PC-4

4. Access PC-4 and run a continuous ping to **8.8.8.8**. Ping should be successful.

Next you will simulate a failure by rebooting the Master unit.

Access-VSF: Member 1

5. Move to Member 1.
6. Reboot it.

```
T11-Access-VSF# vsf member 1 reboot
The master switch will reboot and the standby will become the master.
Do you want to continue (y/n)? y
T11-Access-VSF#
```

PC-4

7. Move back to PC-4.

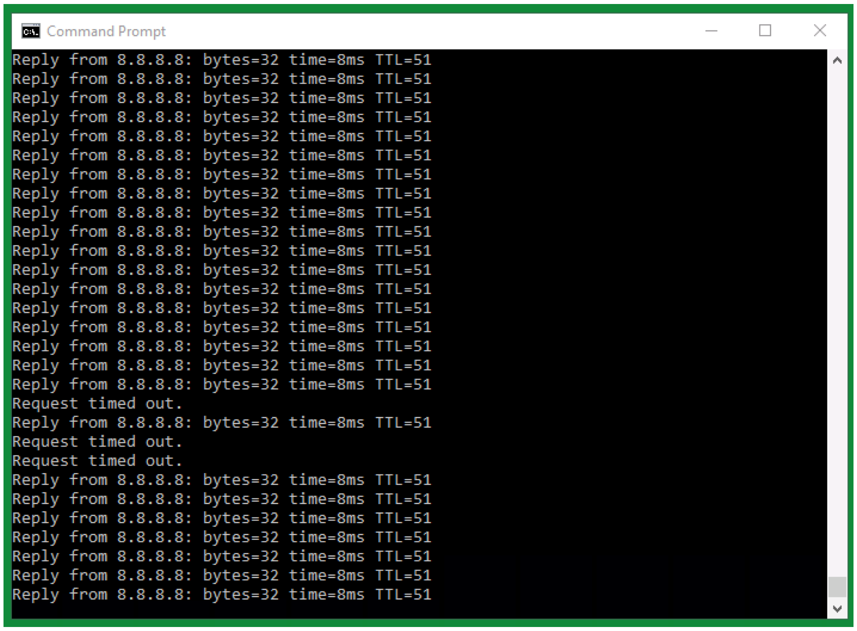

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Request timed out.
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Request timed out.
Request timed out.
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
```

Figure 12.2-3: Ping to internet

Is the ping still running?

____

How many packets did you lose?

____

Access-VSF: Member 2

8. Move to Member 2. As you can see the unit is still alive.
9. Issue the **show vsf** command.

```
T11-Access-VSF# show vsf

MAC Address                 : 88:3a:30:98:30:00
Secondary                   : 2
```

```
Topology                     : Standalone
Status                       : Active Fragment
Split Detection Method       : None


Mbr Mac Address          type           Status
ID
--- ------------------- -------------- ---------------
1                        JL668A         Not Present
2   88:3a:30:97:a4:40    JL668A         Master
T11-Access-VSF#
```

What is the topology?

____

What is the status of the fragment?

____

What role does the member have?

____

10. Wait until Member-1 recovers, then repeat step 9.

```
T11-Access-VSF# show vsf

MAC Address                  : 88:3a:30:98:30:00
Secondary                    : 2
Topology                     : Chain
Status                       : No Split
Split Detection Method       : None


Mbr Mac Address          type           Status
ID
--- ------------------- -------------- ---------------
1   88:3a:30:98:30:00    JL668A         Standby
2   88:3a:30:97:a4:40    JL668A         Master
T11-Access-VSF#
```

What role did Member 1 get when it came back?

NOTE: Master role in VSF is not preemptable: current Master remains the master.

11. Issue the **vsf switchover** for restoring the Master role to Member 1.

```
T11-Access-VSF# vsf switchover
This will cause an immediate switchover to the standby
and the master will reboot.
Do you want to continue (y/n)? y
T11-Access-VSF#
 Feb  4 20:25:49 hpe-mgmtmd[2986]: RebootLibPh1: Reboot triggered due to Reboot
requested through database
```

Access-VSF: Member 1

12. Move to Member 1. You will see that due the “switchover” event, any previous console session that Member 1 had was closed and you will have to login again.

```
The current session terminated due to a failover event. Login again to access the
router.


 (C) Copyright 2017-2019 Hewlett Packard Enterprise Development LP

                     RESTRICTED RIGHTS LEGEND
 Confidential computer software. Valid license from Hewlett Packard Enterprise
 Development LP required for possession, use or copying. Consistent with FAR
 12.211 and 12.212, Commercial Computer Software, Computer Software
 Documentation, and Technical Data for Commercial Items are licensed to the
 U.S. Government under vendor's standard commercial license.

We'd like to keep you up to date about:
  * Software feature updates
  * New product announcements
  * Special events
Please register your products now at: https://asp.arubanetworks.com


T11-Access-VSF login:
```

Task 2: Split Brain Detection

Objectives

After a Master failure, the standby member switch or fragment remains alive. This is because the fragment senses when the links to the Master go down and assumes the Master went down as well. However, what would happen if connections between the two devices fail rather than the master switch? You will discover what happens in the next task.

Steps

PC-3 and PC-4

1. Move to PC-3
2. Run 3 continuous pings to: PC-3's gateway (**10.X.11.254**)**,** PC-4 (**10.X.12.104**) **and 8.8.8.8**. Pings should be successful.
3. Move to PC-4
4. Run 3 continuous pings to: PC-4's gateway (**10.X.12.254**), PC-3 (**10.X.11.103**) and **8.8.8.8**.

Access-VSF: Member 1

5. Move to Member 1
6. Disable the physical port of the VSF link. This will trigger a split-brain event.

```
T11-Access-VSF# config t
T11-Access-VSF(config)# interface 1/1/28
T11-Access-VSF(config-if-vsf)# shutdown
This may cause the stack to split.
Do you want to continue (y/n)? y
T11-Access-VSF(config-if-vsf)#
```

PC-3 and PC-4

7. Move to PC-3.

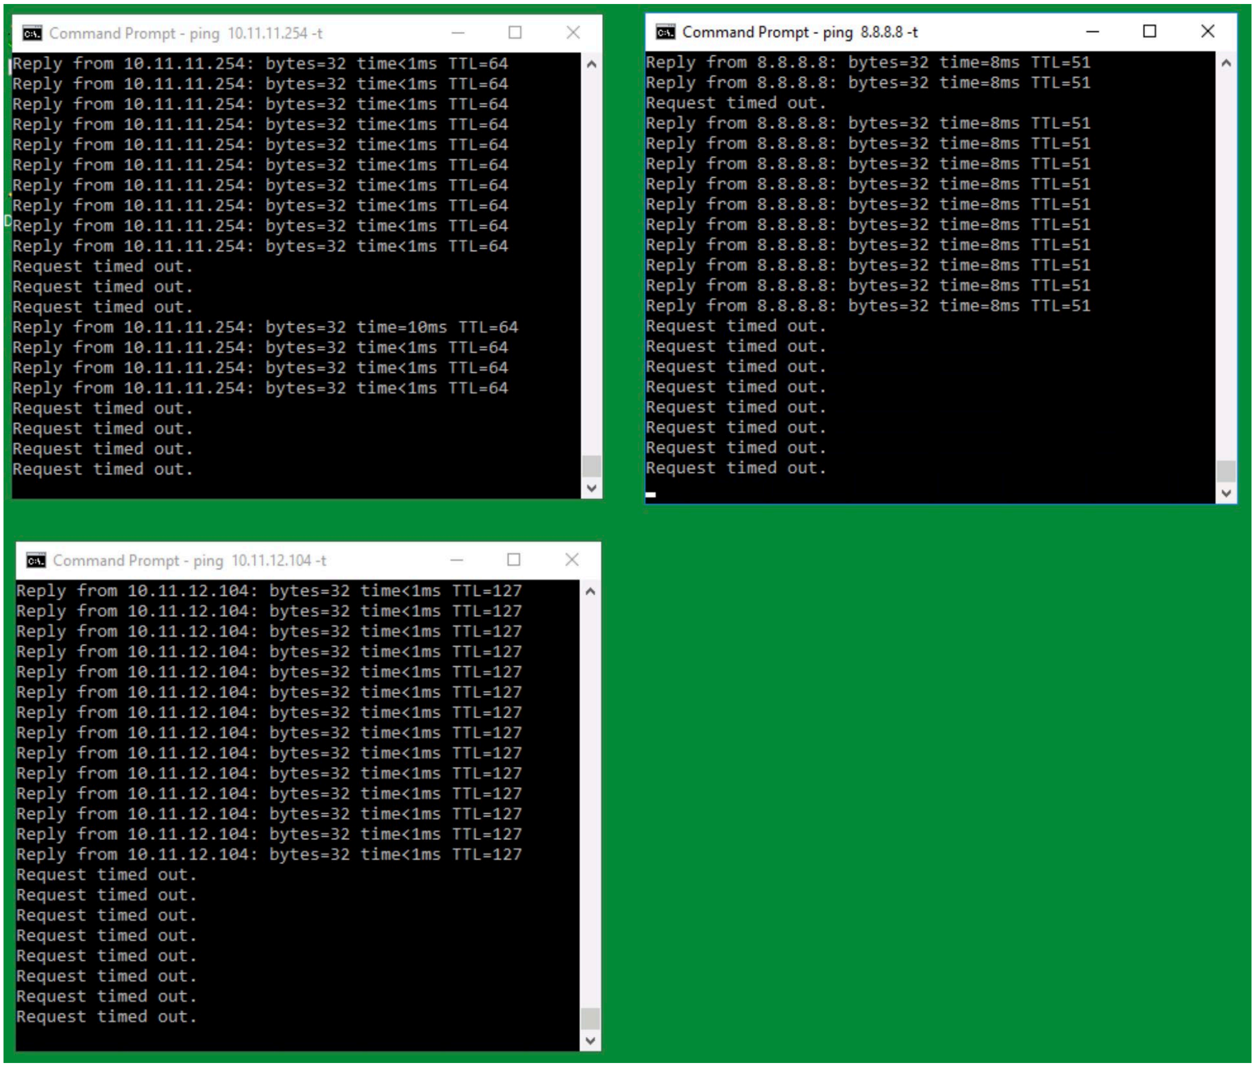

```
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.11.254: bytes=32 time=10ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Request timed out.
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

```
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Reply from 10.11.12.104: bytes=32 time<1ms TTL=127
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

Figure 12.2-4: Multiple pings from PC-3

How are the pings behaving?

____

____

8. Move to PC-4.

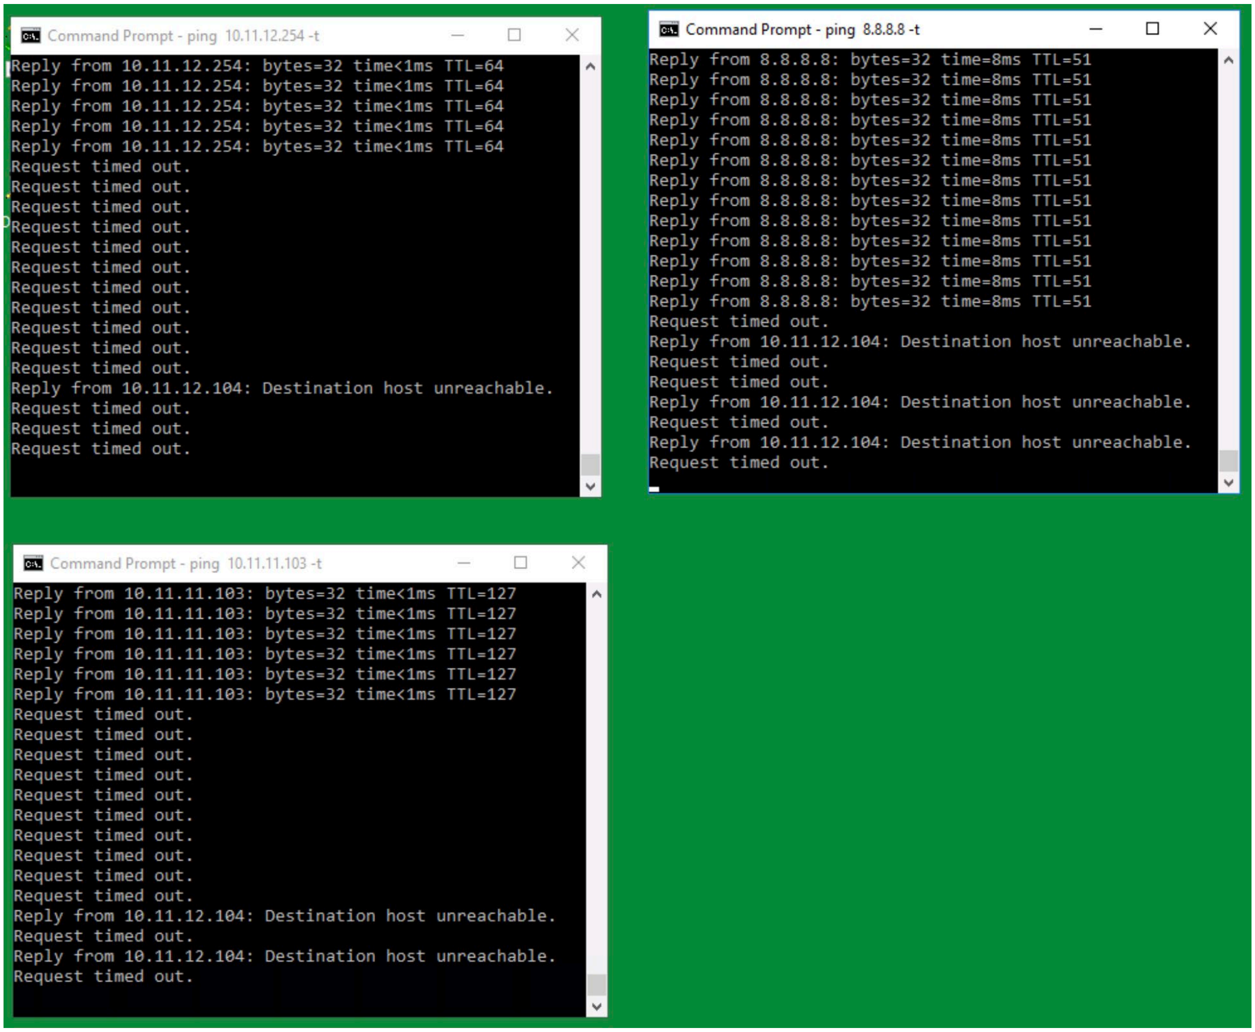

```
Reply from 10.11.12.254: bytes=32 time<1ms TTL=64
Reply from 10.11.12.254: bytes=32 time<1ms TTL=64
Reply from 10.11.12.254: bytes=32 time<1ms TTL=64
Reply from 10.11.12.254: bytes=32 time<1ms TTL=64
Reply from 10.11.12.254: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Request timed out.
```

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
```

```
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Reply from 10.11.11.103: bytes=32 time<1ms TTL=127
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
```

Figure 12.2-5: Multiple pings from PC-4

How are the pings behaving?

____

____

What is the current status of your stack members?

____

Core-1

9. Move **to Core-1.**

10. Issue a filter version of the **show lacp interfaces** command looking for entries containing **lag X1** (this is the lag that connects to your stack).

> **NOTE:** Replace the highlighted "X" for your student table number.

```
Core-1# show lacp interfaces | include X1
1/1/16  lag111     17    1     ASFNCD  90:20:c2:bc:ed:00 65534  111  up
1/1/37  lag111     38    1     ASFNCD  90:20:c2:bc:ed:00 65534  111  up
1/1/16  lag111     22    1     ASFNCD  88:3a:30:98:30:00 65534  111
1/1/37  lag111     86    1     ASFNCD  88:3a:30:98:30:00 65534  111
Core-1#
```

Focus at the first two entries, what is the status of the interfaces?

____

11. Issue the **show interface lag brief** command. The output may be longer than the one below.

```
Core-1# show interface lag brief
--------------------------------------------------------------------------------
----
Port       Native  Mode    Type            Enabled Status   Reason
Speed
           VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
lag10      1       trunk   --              yes     up       --
2000
lag101     1       trunk   --              yes     up       --
2000
lag111     1       trunk   --              yes     up       --
2000
Core-1#
```

> **NOTE:** Since Core-1 is a shared resource you may get more entries in the command's output.

What is the status of lag X1?

IMPORTANT: The problem you are experiencing is a result of having two stack fragments (member 1 and member 2) both acting as masters and using not only the same configuration, but also the same Layer 3 and Layer 2 addressing. Therefore, they are sending identical LACP Data Units on the interfaces that are configured to be part of the same lag (1/1/21 and 2/1/21 to Core-1 and 1/1/22 and 2/1/22 to Core-2). See figure 12.2-6.

Since the Core switches receive these incoming LACP Data Units as normal, they are not aware of any failure and maintain their LAGs and forward traffic across them as usual, based on the source and destination IP addresses.

```
Core-1# show lacp aggregates lagX1

Aggregate name   : lag111
Interfaces       : 1/1/16 1/1/37
Heartbeat rate   : Fast
Hash             : l3-src-dst
Aggregate mode   : Active

Core-1#
```

Depending on what hash result the Core switches calculate for each of the pings, the traffic path could be similar to the one shown in figure 12.2-7. In this example, PC-1 tries to reach 8.8.8.8 and delivers the packets to Member 1. Member 1 then forwards the packets across the local port member of Lag X1 (1/1/21). Core-1 gets the packet and just routes it as normal. There are no abnormal behaviors yet.

The problem arises when Core-1 receives the reply from 8.8.8.8. Instead of sending it straight to Member 1 using interface 1/1/Y, it gives it to Member 2 via port 1/1/Z as result of the hashing algorithm, then Member 2, not having any physical connection to PC-1, has no option but to drop the packet.

Although some traffic flows might work, many others will not. The unpredictable nature of this outcome makes the network unusable when split-brain takes place.

NOTE: If your connectivity test from PC-3 to 8.8.8.8 are still working successfully, then it is likely that the behavior explained in the lines above is taking place on another of your pings.

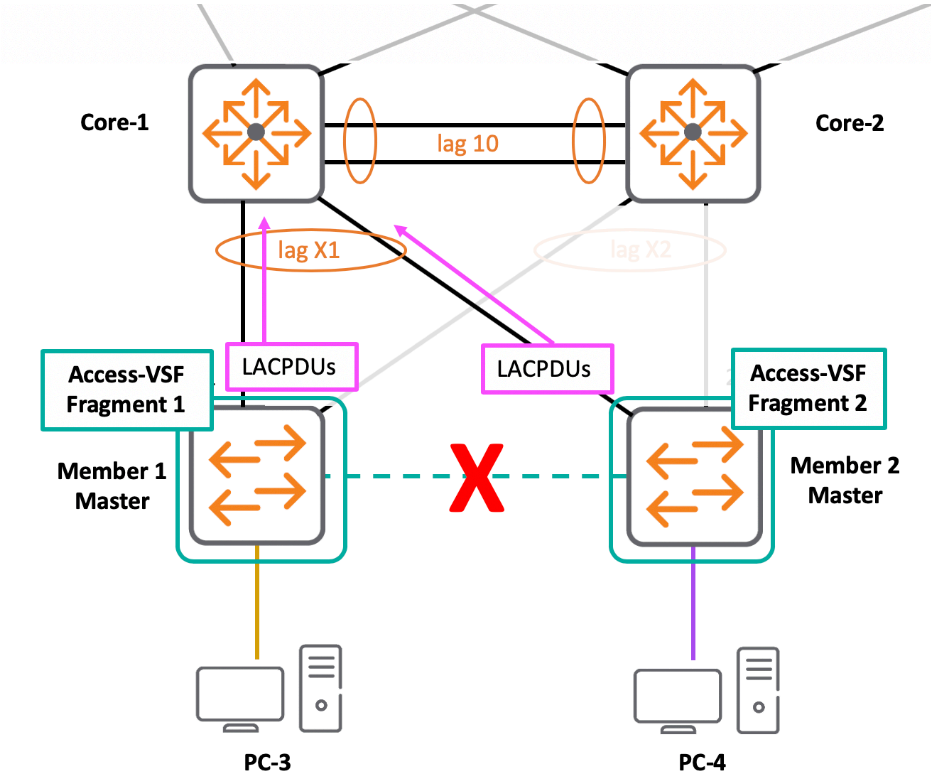

Figure 12.2-6: Split Brain

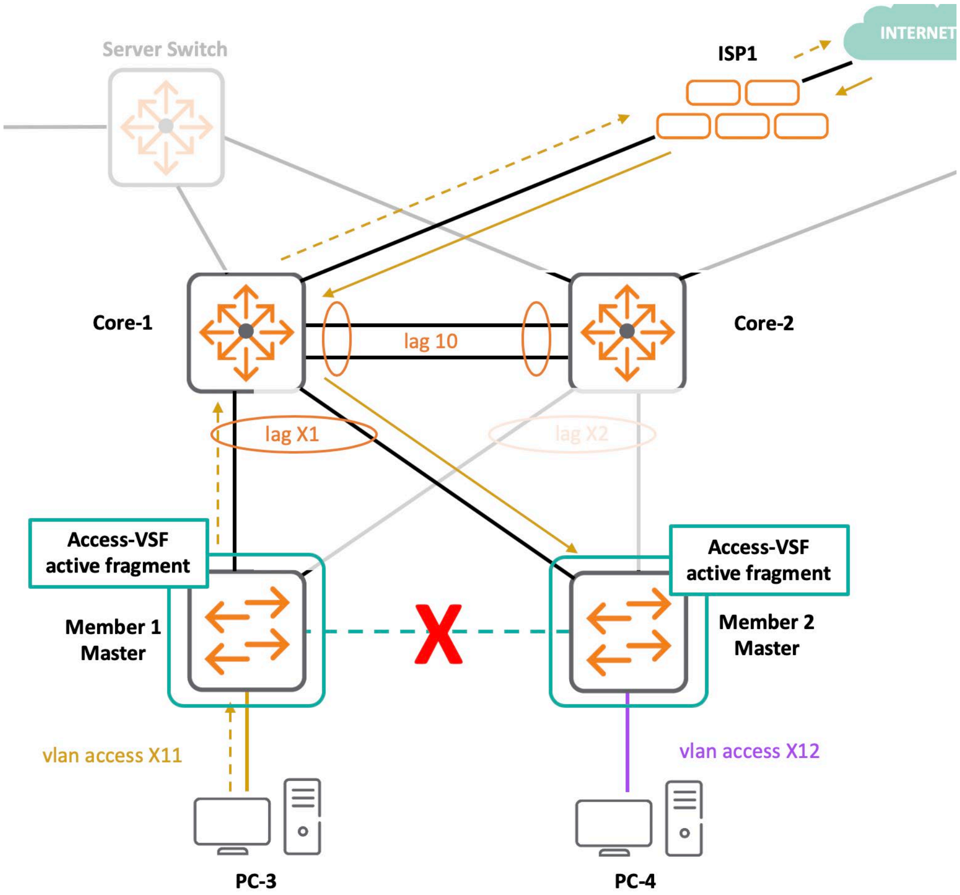

Figure 12.2-7: Split Brain consequences

Access-VSF: Member 1

12. Move back to Member-1.
13. Enable the port of the VSF link. Member-2 will merge and reboot.

```
T11-Access-VSF(config-if-vsf)# no shutdown
T11-Access-VSF(config-if-vsf)#
```

Access-VSF: Member 2

14. Move to Member-2. You shall notice the member switch will reboot as part of the re-merge process.

```
T11-Access-VSF#
 Feb  4 20:48:14 vsfd[719]: RebootLibPh1: Reboot triggered due to Reboot of
Member ID 2, Lost merge
```

NOTE: Now you will enable management port based, split brain detection. When this feature is enabled, the Master and Standby Member will exchange broadcast-based heartbeats when they sense a failure in the VSF links. If the Standby member does not receive any of these messages, then it concludes that the Master itself has failed, not just the VSF links. Therefore, it keeps working as normal. However, if the Master is alive and continues to advertise split detect messages, then the Standby Member's fragment changes its status to inactive and disables all its ports except the management and VSF interfaces. This isolates it from the rest of the network and prevents the Cores from sending traffic to it.

Although this behavior will affect every endpoint connected to the inactive fragment, those connected to the Active one will not have any connection loss and will always be able to establish connections with any destination in the network, with the exception of clients connected directly to the inactive fragment.

Access-VSF: Member 1

15. Move back to member 1.
16. Enable **split-detection**.

```
T11-Access-VSF(config)# vsf split-detect mgmt
T11-Access-VSF(config)#
```

17. Issue the **show vsf** command and confirm Split Detection Method is **mgmt**.

```
T11-Access-VSF(config)# show vsf

MAC Address                : 88:3a:30:98:30:00
Secondary                  : 2
Topology                   : Chain
Status                     : No Split
Split Detection Method     : mgmt


Mbr Mac Address         type           Status
```

```
ID
--- ------------------- -------------- ---------------
1   88:3a:30:98:30:00   JL668A         Master
2   88:3a:30:97:a4:40   JL668A         Standby
T11-Access-VSF(config)#
```

18. Disable the physical port of the VSF link. This will trigger split-detect messages from the Standby Member, see figure 12.2-8.

```
T11-Access-VSF(config)# interface 1/1/28
T11-Access-VSF(config-if-vsf)# shutdown
This may cause the stack to split.
Do you want to continue (y/n)? y
T11-Access-VSF(config-if-vsf)#
```

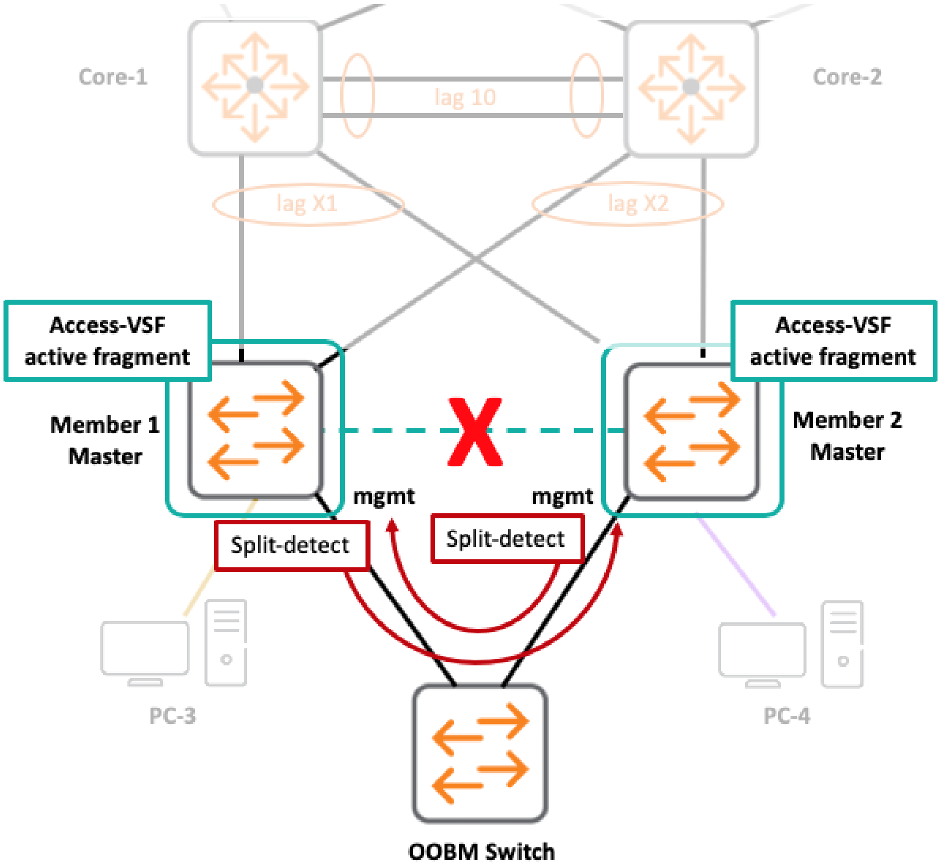

Figure 12.2-8: Split detect data units

NOTICE: Split detect uses ethertype **0xf8f8**, if you happen to deploy any layer 2 filtering tool in the Out of Band Management switch, then make sure these packets are explicitly permitted.

PC-3 and PC-4

19. Move back to PC-3.

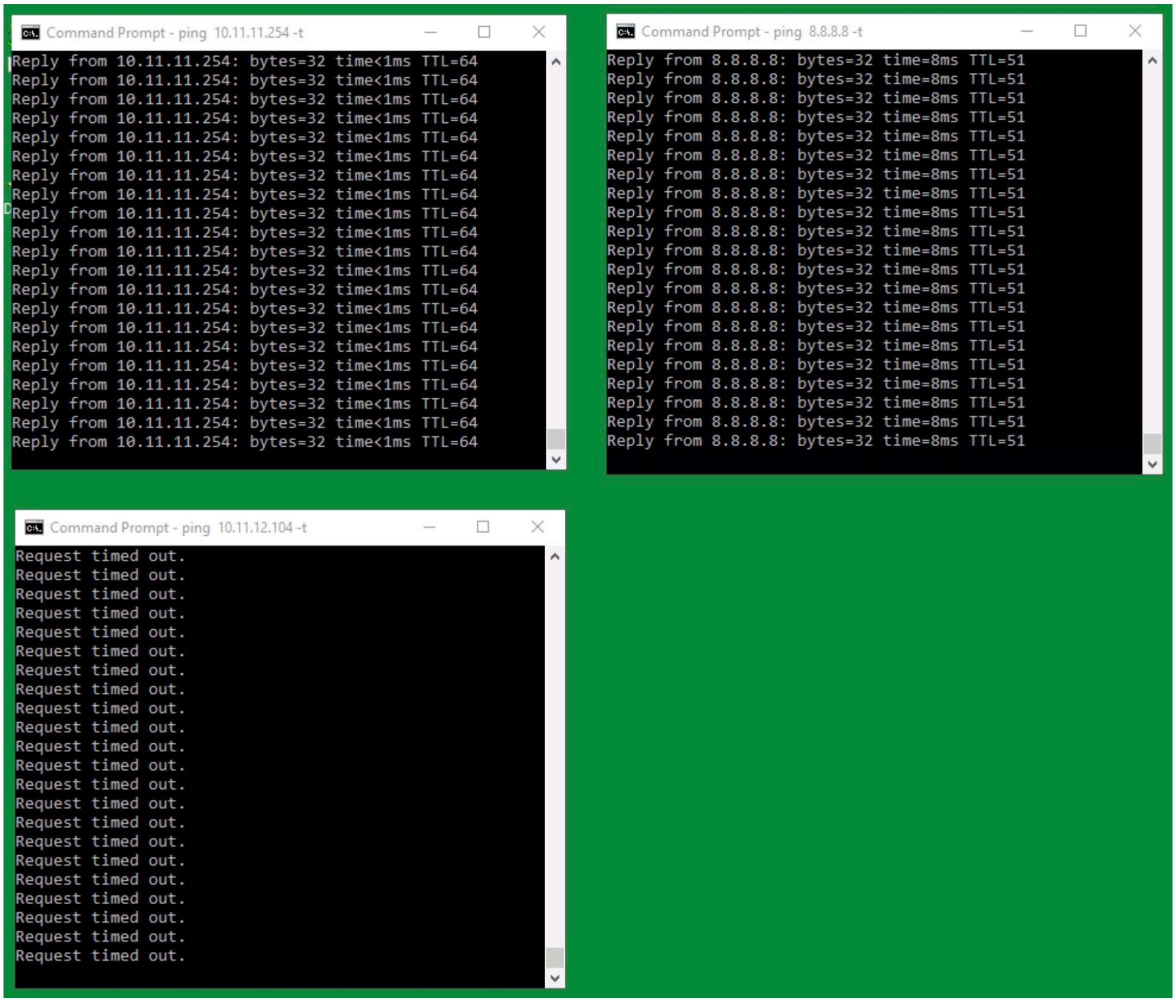

```
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
Reply from 10.11.11.254: bytes=32 time<1ms TTL=64
```

```
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
Reply from 8.8.8.8: bytes=32 time=8ms TTL=51
```

```
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

Figure 12.2-9: Multiple pings from PC-3 2

Are pings still running?

____

Which one is falling?

Is this result what you expected?

---

20. Move back to PC-4.

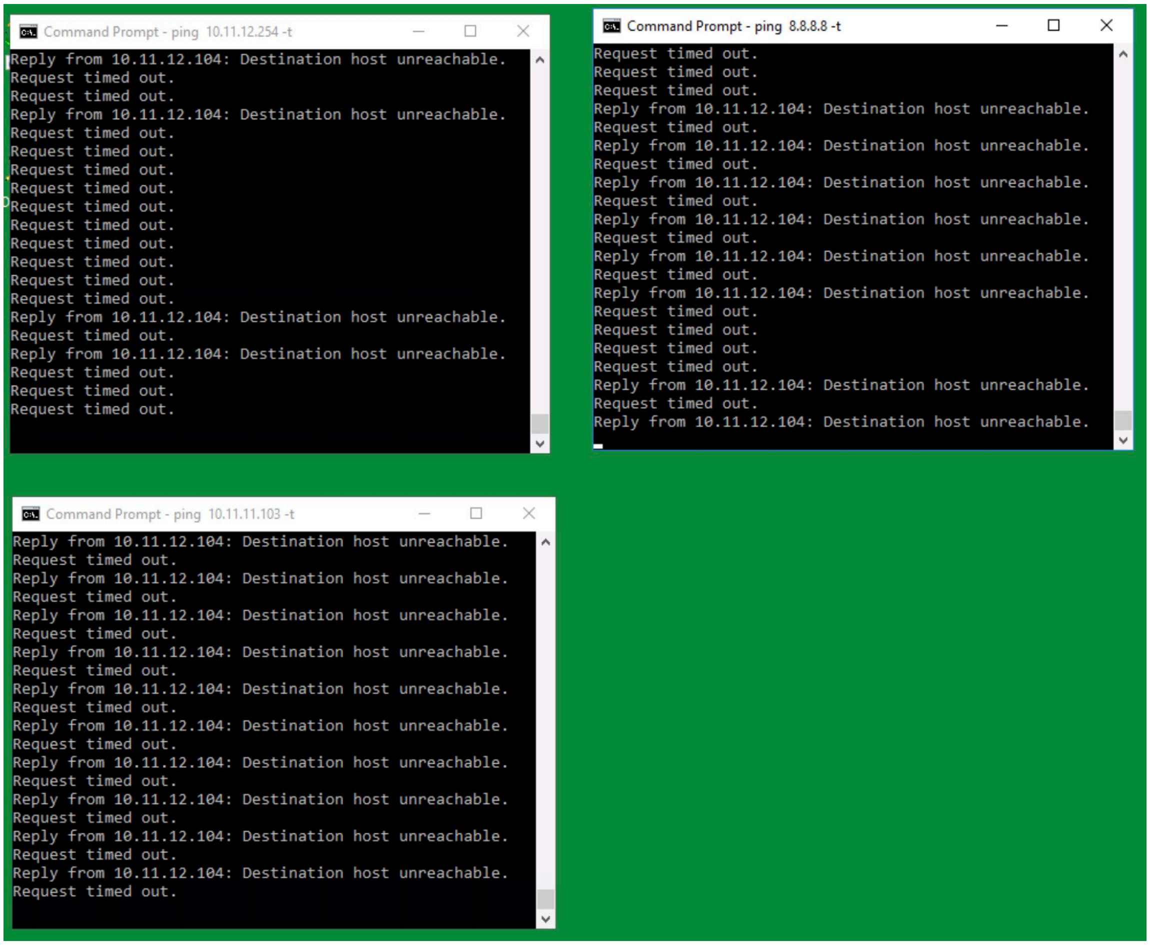

```
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Request timed out.
```

```
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
```

```
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
Reply from 10.11.12.104: Destination host unreachable.
Request timed out.
```

Figure 12.2-10: Multiple pings from PC-4 2

Are pings still running?

Which one is falling?

Is this result what you expected?

Access-VSF: Member 1

21. Move back to Member 1.
22. Issue the **show vsf** command.

```
T11-Access-VSF(config-if-vsf)# show vsf

MAC Address                : 88:3a:30:98:30:00
Secondary                  : 2
Topology                   : Standalone
Status                     : Active Fragment
Split Detection Method     : mgmt


Mbr Mac Address          type           Status
ID
--- ------------------- -------------- ---------------
1   88:3a:30:98:30:00   JL668A         Master
2                       JL668A         In Other Fragment
T11-Access-VSF(config-if-vsf)#
```

What is the status of the fragment?

What is the status of Member 2?

Access-VSF: Member 2

23. Move back to Member 2.
24. Repeat step 21.

```
T11-Access-VSF# show vsf

MAC Address                : 88:3a:30:98:30:00
Secondary                  : 2
Topology                   : Standalone
Status                     : Inactive Fragment
Split Detection Method     : mgmt


Mbr Mac Address           type           Status
ID
--- ------------------- -------------- ---------------
1                         JL668A         In Other Fragment
2   88:3a:30:97:a4:40     JL668A         Master
T11-Access-VSF#
```

What is the status of the fragment?

____

What is the status of Member 2?

____

25. Use the **show interface brief** command and look for the status of both uplinks and connection to PC-4.

```
T11-Access-VSF# show interface brief | exclude no
--------------------------------------------------------------------------------
Port       Native  Mode   Type           Enabled Status  Reason
Speed
           VLAN
(Mb/s)
---------------------------------------------------------------------------------
2/1/4      1112    access 1GbT           yes     down    Disabled by VSF        --
2/1/21     1       trunk  1GbT           yes     down    Disabled by VSF        --
2/1/22     1       trunk  1GbT           yes     down    Disabled by VSF        --
2/1/28     --      routed SFP+DAC1       yes     down    Waiting for link       --
```

```
vlan1     --            --            yes     down                               --
lag111    1      trunk  --            yes     down    --                 auto
lag112    1      trunk  --            yes     down    --                 auto
T11-Access-VSF#
```

What is the status of these ports?

What is the reason?

Your results should be similar to the one shown in figure 12.2-11.

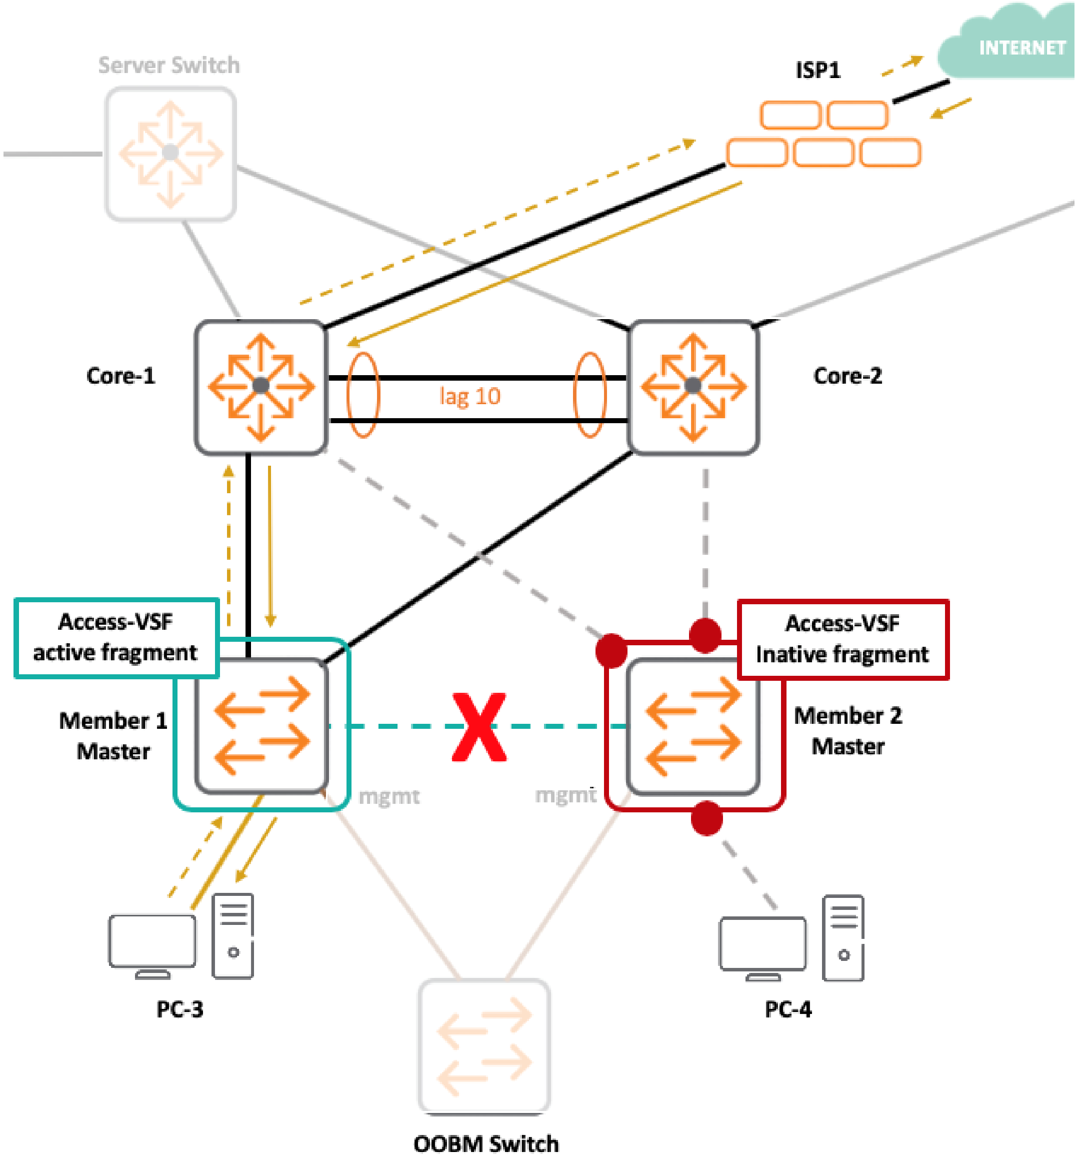

Figure 12.2-11: Inactive fragment

Access-VSF: Member 1

26. Move back one last time to Member 1.
27. Enable the ports.

```
T11-Access-VSF(config-if-vsf)# no shutdown
T11-Access-VSF(config-if-vsf)# end
```

Task 3: Save Your Configurations

Objectives

You will now proceed to save your configurations and create checkpoints. Notice that final lab checkpoints might be used by later activities.

Steps

Access-VSF, Core-1 and Core-2 (via PC-1).

1. Save the current Access-VSF's configuration in the startup checkpoint.

```
T11-Access-VSF # write memory
Configuration changes will take time to process, please be patient.
T11-Access-VSF
 #
Access-2#
```

Access-VSF

2. Backup the current Access-VSF's configuration as a custom checkpoint called **Lab12-2_final**.

```
T11-Access-VSF # copy running-config checkpoint Lab12-2_final
Configuration changes will take time to process, please be patient.
T11-Access-VSF#
```

You have completed Lab 12.2!

AOS-CX Switching Fundamentals

Lab 13: Secure Management Access

Overview

After deploying VSF and instructing the staff member how to gain console access to the system, you get a few queries from him and his manager. They commented that going to the IDF every time a change is needed, consumes a considerable amount of time. They ask if remote access is possible, since they have it with the Core switches. Additionally, they are also interested in any graphical interface alternatives for monitoring system parameters like, CPU, memory, ports and the stack status.

After the meeting, the manager comments behind closed doors that he is aware of the technician's limited training. He wants to restrict the technician's configuration tasks to provisioning the first nine ports of each stack member into the proper VLAN.

Objectives

After completing this lab, you will be able to:

- Enable remote access in the through the mgmt. port only
- Enable local command authorization
- Deploy RADIUS based AAA Role Based Access Control
- Explore AOS-CX web-based UI

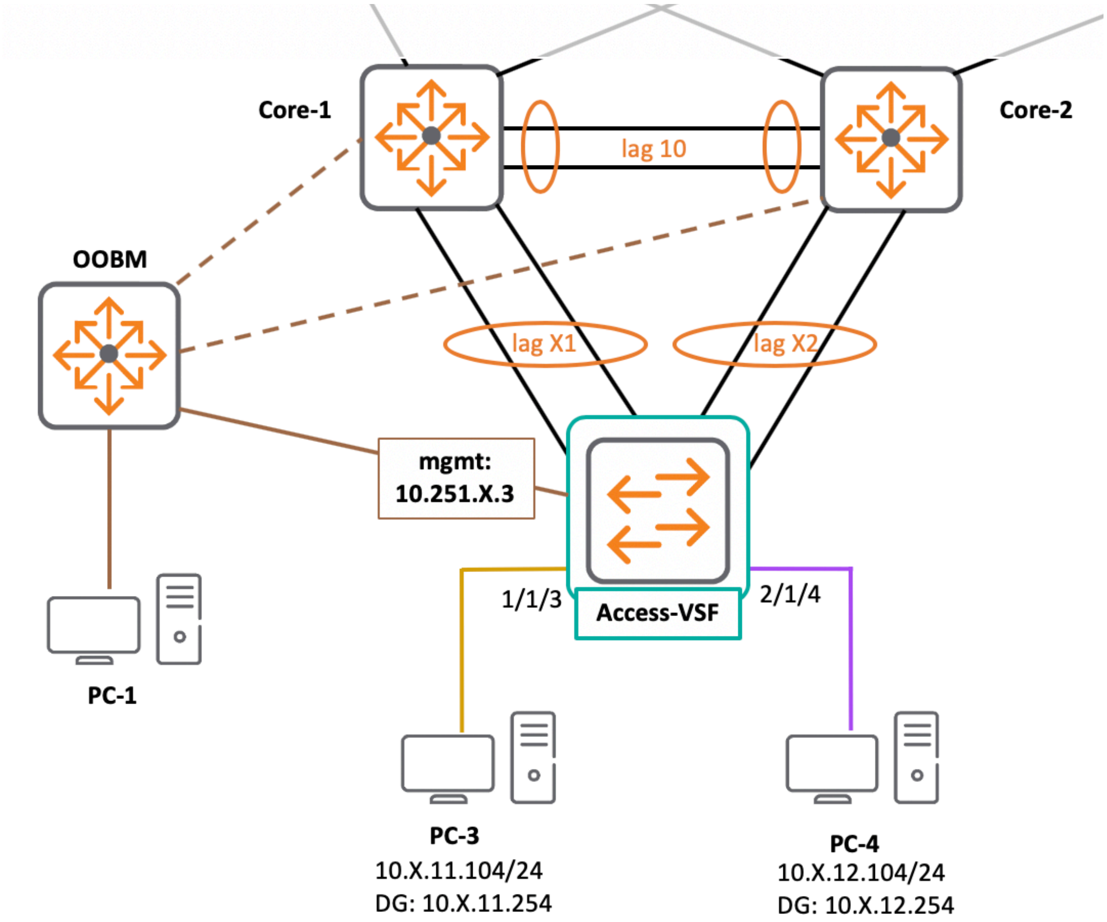

Figure 13-1: Lab Topology

Task 1: Management Port

Objectives

In order to comply with your customer's requirements, you first must assign an IP address to the management port. Remember, this port belongs to an exclusive management specific VRF. Unlike regular data VRFs, where either static or dynamic routing is supported, the management one uses a default gateway, as if it was a host.

Steps

Access-VSF: Member 1

1. Access Member 1's console session.
2. Move to **mgmt** interface and assign the **10.251.X.3/24** IP address.

> **NOTE:** Replace the highlighted "X" for your student table number.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface mgmt
T11-Access-VSF(config-if-mgmt)# ip static 10.251.X.3/24
```

3. Assign the **10.251.X.254** and **10.254.1.22** as gateway and DNS servers respectively.

```
T11-Access-VSF(config-if-mgmt)# default-gateway 10.251.X.254
T11-Access-VSF(config-if-mgmt)# nameserver 10.254.1.22
T11-Access-VSF(config-if-mgmt)# exit
```

4. Display the "mgmt" VRF.

```
T11-Access-VSF(config)# show vrf mgmt
VRF Configuration:
------------------
VRF Name    : mgmt
use "show interface mgmt" for mgmt interfaces
T11-Access-VSF(config)#
```

What interfaces are associated to this VRF?

____

5. Display the “mgmt” interface settings. Confirm the parameters are correct.

```
T11-Access-VSF(config)# show interface mgmt
  Address Mode                    : static
  Admin State                     : up
  Mac Address                     : 88:3a:30:98:30:01
  IPv4 address/subnet-mask        : 10.251.11.3/24
  Default gateway IPv4            : 10.251.11.254
  IPv6 address/prefix             :
  IPv6 link local address/prefix:
  Default gateway IPv6            :
  Primary Nameserver              : 10.254.1.22
  Secondary Nameserver            :
T11-Access-VSF(config)#
```

6. Ping the default gateway (**10.251.X.254**). Ping should be successful.

```
T11-Access-VSF(config)# do ping 10.251.X.254 vrf mgmt
PING 10.251.11.254 (10.251.11.254) 100(128) bytes of data.
108 bytes from 10.251.11.254: icmp_seq=1 ttl=64 time=4.20 ms
108 bytes from 10.251.11.254: icmp_seq=2 ttl=64 time=0.224 ms
108 bytes from 10.251.11.254: icmp_seq=3 ttl=64 time=0.231 ms
108 bytes from 10.251.11.254: icmp_seq=4 ttl=64 time=0.230 ms
108 bytes from 10.251.11.254: icmp_seq=5 ttl=64 time=0.224 ms

--- 10.251.11.254 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4076ms
rtt min/avg/max/mdev = 0.224/1.022/4.201/1.589 ms
```

7. Display the SSH servers on all VRFs.

```
T11-Access-VSF(config)# show ssh server all-vrfs | include SSH
SSH server configuration on VRF default :
    IP Version          : IPv4 and IPv6        SSH Version             : 2.0
SSH server configuration on VRF mgmt :
    IP Version          : IPv4 and IPv6        SSH Version             : 2.0
T11-Access-VSF(config)#
```

What VRFs have an SSH server?

8. Display the SSH servers on all VRFs.

```
T11-Access-VSF(config)# show https-server

HTTPS Server Configuration
----------------------------
 VRF                  : mgmt, default

 REST Access Mode   : read-write

T11-Access-VSF(config)#
```

What VRFs have an HTTPS server?

What VRFs have an HTTPS server?

NOTE: In 6300 and 6400 series switches, SSH and HTTPS services are running by default in both the “mgmt” and “default” VRFs, however in the case of 8300 and 8400s these services are only running in the “mgmt” VRF.

Also REST Access mode comes as read-write in the 6000 platforms, while in the 8000s it begins as “read-only”.

9. Disable SSH and HTTPS services from default VRF. This will prevent this traffic from being processed in the regular data VRF.

```
T11-Access-VSF(config)# no ssh server vrf default
Active SSH sessions will be terminated.
Do you want to continue (y/n)? y
T11-Access-VSF(config)# no https-server vrf default
T11-Access-VSF(config)#
```

Task 2: RBAC

Objectives

The next step to comply with your customer's desires is to enable local command authorization. That is achieved by creating user groups and local user accounts in AOS-CX. In this task, you are going to define a list of allowed commands. In order to reduce the number of lines needed for the task you will leverage the power of Regular Expressions (REGEX).

> **NOTICE:** Regular expressions are text strings used for describing a search pattern, they are considered the next step in the evolution of wildcards. Several features and tools in networking, IT, engineering, science, etc. use REGEX for matching strings. You might find it useful to start learning about them.

Steps

Access-VSF: Member 1

1. Access Member 1's console session.
2. Create a user group called "**port-prov**", then allow the following:
    a. Access to global configuration context.
    b. Access to first 9 ports on both VSF members.
    c. Change VLAN membership on those ports.
    d. Enable ports.
    e. Display a list of interfaces, VLANs, and user information.

```
T11-Access-VSF(config)# user-group port-prov
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "configure terminal"
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "interface [1-
2]/1/[1-9]$"
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "vlan access"
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "no shutdown"
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "show interface
brief"
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "show vlan"
```

```
T11-Access-VSF(config-usr-grp-port-prov)# permit cli command "show user
information"
T11-Access-VSF(config-usr-grp-port-prov)# exit
```

TIP: Defining commands of different user groups supports REGEX. For example, in the second rule [1-2] means that the character could take either value “1” or “2”, likewise [1-9] represents any number in the range between 1 to 9, and “$” means this is the end of the line and nothing else can follow.

3. Display the user-groups.

```
T11-Access-VSF(config)# show user-group
GROUP NAME     GROUP TYPE     INCLUDED GROUP     NUMBER OF RULES
-------------- -------------- ------------------ -------------------
administrators built-in       n/a                n/a
auditors       built-in       n/a                n/a
operators      built-in       n/a                n/a
port-prov      configuration  --                 7
T11-Access-VSF(config)#
```

In addition to “**port-prov**” what groups are listed?

____

NOTE: The “operator” context enables you to execute commands to view, but not change, the configuration. This group has privilege level 1.

Users with “auditor” rights have access to show accounting, events and logging commands and the ability to use copy show commands to direct output onto a USB or remote storage. The prompt for this kind of session is “auditor>”. This group has privilege level 19.

Administrator group grants “manager” access (full access) to every aspect of the system. This group has privilege level 15.

4. Display the details of your group. You will notice all the rules you have defined with sequence numbers in steps of 10.

```
T11-Access-VSF(config)# show user-group port-prov
User Group Summary
==================
Name             : port-prov
Type             : configuration
Included Group   : --
Number of Rules  : 7

User Group Rules
================
SEQUENCE NUM  ACTION     COMMAND                         COMMENT
------------- ---------- ------------------------------- -----------------------
---------
10            permit     configure terminal
20            permit     interface [1-2]/1/[1-9]$
30            permit     vlan access
40            permit     no shutdown
50            permit     show interface brief
60            permit     show vlan
70            permit     show user information
T11-Access-VSF(config)#
```

5. Create the **cxfX-local** user account password **aruba123**. Map the account to the **port-prov** group you just created.

NOTE: Replace the highlighted “X” for your student table number.

```
T11-Access-VSF(config)# user cxfX-local group port-prov password plaintext
aruba123
T11-Access-VSF(config)#
```

6. Display the local user list. You will see only two accounts.

```
T11-Access-VSF(config)#  show user-list
USER                             GROUP
---------------------------------------
admin                            administrators
cxf11-local                      port-prov
T11-Access-VSF#
```

NOTE: Although the scenario is asking for secure RBAC, the “**admin**” account should remain untouched with no password. This eases the assistance and reset procedures that the lab help desk might need to run.

PC-1

7. Access PC-1's console session.
8. Open putty.
9. Run an SSH session to the management IP address of the Access-VSF (**10.251.X.3**).

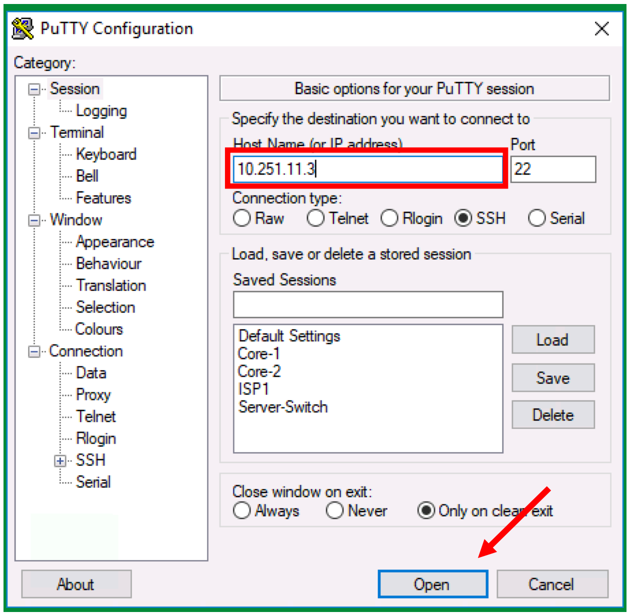

Specify the destination you want to connect to

Host Name (or IP address): 10.251.11.3

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Saved Sessions: Default Settings, Core-1, Core-2, ISP1, Server-Switch

Close window on exit: ○ Always ○ Never ● Only on clean exit

Figure 13-2:Putty

Access-VSF (via PC-1)

10. Login with **cxfX-local/aruba123.**

```
login as: cxfX-local

 (C) Copyright 2017-2019 Hewlett Packard Enterprise Development LP

                         ←---- output omitted ---→

cxf11@10.251.11.3's password: aruba123

 T11-Access-VSF#
```

11. Try the **show user information** command. You shall see the user you are using for this session and the user group it belongs.

```
T11-Access-VSF# show user information
Username              : cxf11-local
Authentication type   : local
User group            : port-prov
User privilege level : N/A
T11-Access-VSF#
```

12. Move port **2/1/4** to **VLAN X11**.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface 2/1/4
T11-Access-VSF(config-if)# vlan access X11
T11-Access-VSF(config-if)# end
```

13. Display **VLAN X11** and confirm port **2/1/4** is there.

```
T11-Access-VSF# show vlan X11

---------------------------------------------------------------------------------
-----------------------------
VLAN  Name                                  Status  Reason                Type
Interfaces
---------------------------------------------------------------------------------
-----------------------------
1111  EMPLOYEES                             up      ok                    static
1/1/3,2/1/4,lag111-lag112
T11-Access-VSF#
```

14. Display the running configuration.

```
T11-Access-VSF# show running-config
Cannot execute command. Command not allowed.
T11-Access-VSF#
```

15. Access **lag X1** interface then port **1/1/10**.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface lag X1
Cannot execute command. Command not allowed.
```

```
T11-Access-VSF(config)#
```

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface 1/1/10
Cannot execute command. Command not allowed.
T11-Access-VSF(config)#
```

Did you experience any issues trying out any of those 3 commands?

____

What is most likely the cause?

____

____

PC-4

16. Move to PC-4.
17. Run Command Prompt as administrator.

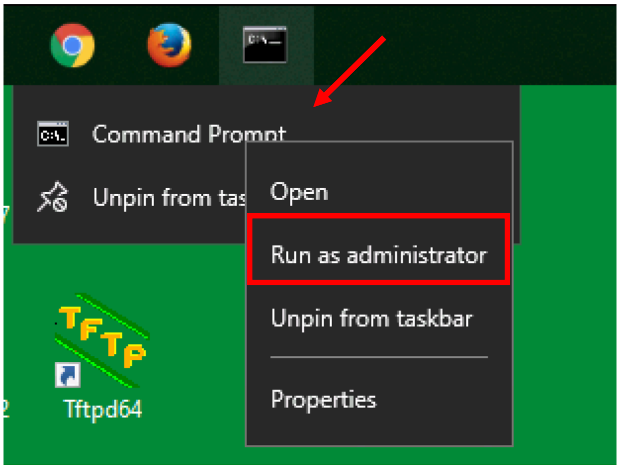

Figure 13-3:Run as administrator

18. Run **ipconfig -renew** to request an IP address from **VLAN X11**.

TIP: If you are not allowed to run the command then make sure your NIC is setup as DHCP client.

Task 3: RADIUS Based Management

Objectives

After testing the command-based authorization with your customer and demonstrating the power of this management control, you explain that local accounts are not always the best option, especially with fast growing networks like BigStartup. Having operator accounts that need to be created at every single switch and system does not scale well - especially when a password change or account revocation is required. Therefore, you offer to deploy a ClearPass demo in order to give them a taste of account centralization and also demonstrate the power of the ClearPass product.

In this task you will enable RADIUS based authentication for SSH and HTTPS sessions.

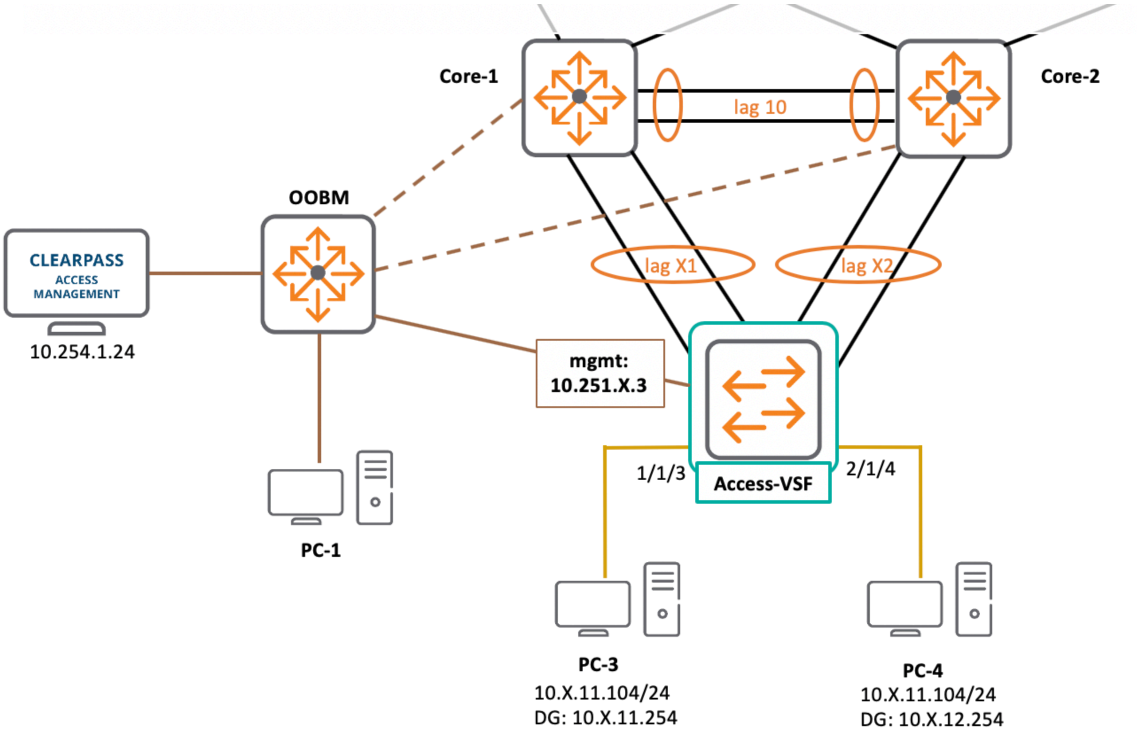

Figure 13-4: OOBM Network

Steps

Access-VSF: Member 1

1. Access Member 1’s console session.
2. Define a RADIUS server with the **10.254.1.24** IP address. Use **aruba123** as the shared secret and map it to **VRF mgmt**.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# radius-server host 10.254.1.24 key plaintext aruba123 vrf
mgmt
T11-Access-VSF(config)#
```

3. Display the newly created RADIUS server. Confirm all settings are in order.

```
T11-Access-VSF(config)# show radius-server | begin 3 SERVER

SERVER NAME                                   | PORT | VRF
---------------------------------------------------------------------------------
10.254.1.24                                   | 1812 | mgmt
---------------------------------------------------------------------------------
T11-Access-VSF(config)#
```

4. Set the RADIUS group then the local username database as authentication groups for HTTPS and SSH services.

```
T11-Access-VSF(config)# aaa authentication login https-server group radius local
T11-Access-VSF(config)# aaa authentication login ssh group radius local
T11-Access-VSF(config)#
```

NOTICE: It is best practice to have a local database backup a remote authentication group when configuring AAA management access. This prevents locking out the administrator’s account if the AAA server fails or becomes unreachable.

PC-1

5. Access PC-1’s console session.
6. Run an SSH session to the management IP address of the Access-VSF.

```
login as: cxfX

 (C) Copyright 2017-2019 Hewlett Packard Enterprise Development LP

                         ←---- output omitted ---→

cxf11@10.251.11.3's password: aruba123

T11-Access-VSF#
```

7. Login with **cxfX/aruba123**. This account is stored in ClearPass.
8. Try the **show user information** command.

```
T11-Access-VSF# show user information
Username             : cxf11
Authentication type  : RADIUS
User group           : administrators
User privilege level : 15
T11-Access-VSF#
```

What is the Authentication type?

To what user group does the user belong?

What is the privilege level?

You will now proceed to remove local command authorization.

9. Delete the local account.

```
T11-Access-VSF(config)# no user cxfX-local
User cxf11-local's home directory and active sessions will be deleted.
Do you want to continue (y/n)? y
T11-Access-VSF(config)#
```

10. Delete the port-prov authorization group.

```
T11-Access-VSF(config)# no user-group port-prov
T11-Access-VSF(config)#
```

Task 4: Explore AOS-CX Web UI

Objectives

Your customer's final request is to use a graphical interface for monitoring the system. You invite the executives from BigStartup to explore AOS-CX Web User Interface.

Steps

PC-1

1. Access the console session of PC-1.
2. Open a browser and point it to your Access-VSF management IP address (**10.251.X.3**).

> **NOTE:** Replace the highlighted "X" for your student table number.

3. Login using **cxfX/aruba123**

> **NOTE:** Replace the highlighted "X" for your student table number, e.g. username: **cxf4**, password: **aruba123**. Or username: **cxf11**, password: **aruba123,** for tables 4 and 11 respectively.

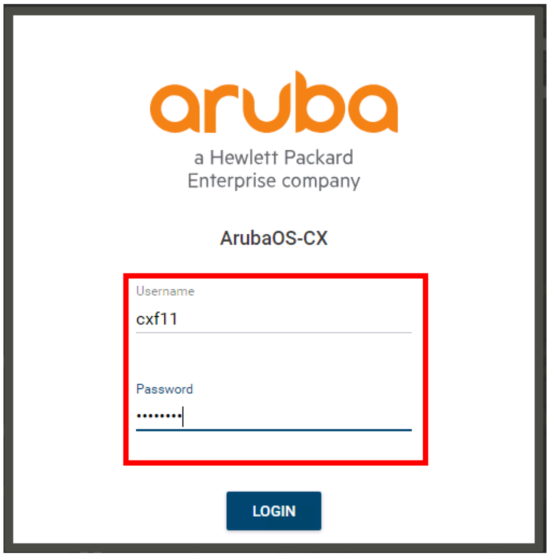

Username
cxf11

Password
••••••••

Figure 13-5: Web login page AOS-CX

4. Accept the pre-login banner. You will be taken to the **Overview** page.

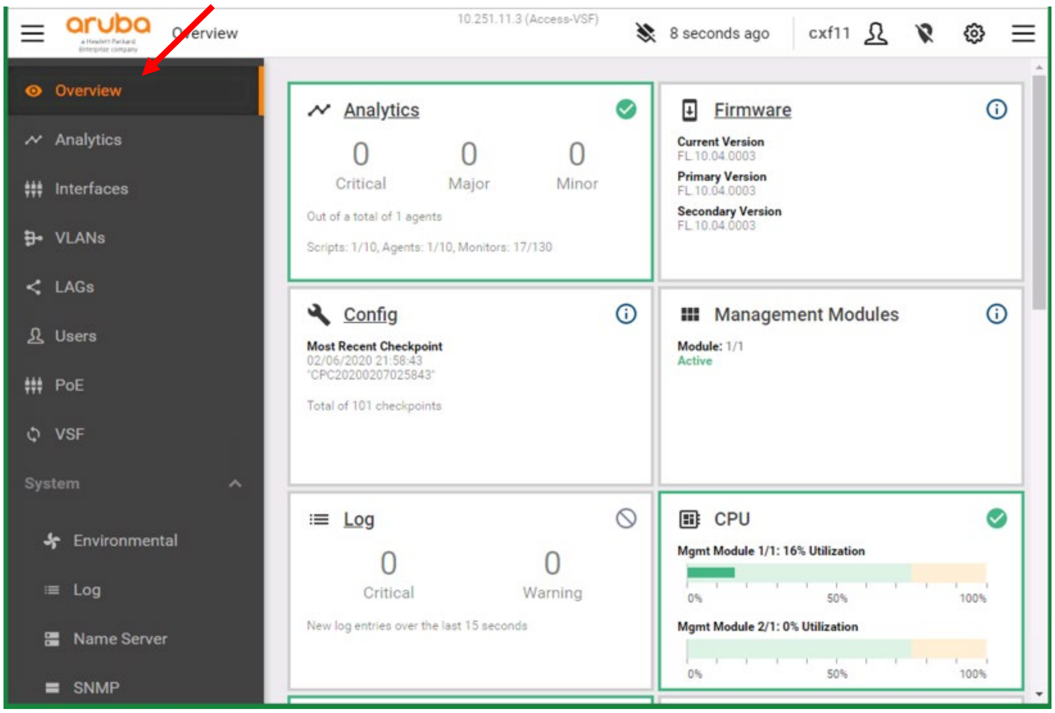

Analytics: 0 Critical, 0 Major, 0 Minor
Out of a total of 1 agents
Scripts: 1/10, Agents: 1/10, Monitors: 17/130

Firmware
Current Version
FL.10.04.0003
Primary Version
FL.10.04.0003
Secondary Version
FL.10.04.0003

Config
Most Recent Checkpoint
02/06/2020 21:58:43
"CPC20200207025843"
Total of 101 checkpoints

Management Modules
Module: 1/1
Active

Log: 0 Critical, 0 Warning
New log entries over the last 15 seconds

CPU
Mgmt Module 1/1: 16% Utilization
Mgmt Module 2/1: 0% Utilization

Figure 13-6: Overview

What is the Firmware version running in the stack?

____

Are there any new logs?

____

What is the CPU utilization on each stack member?

____

5. Scroll down.

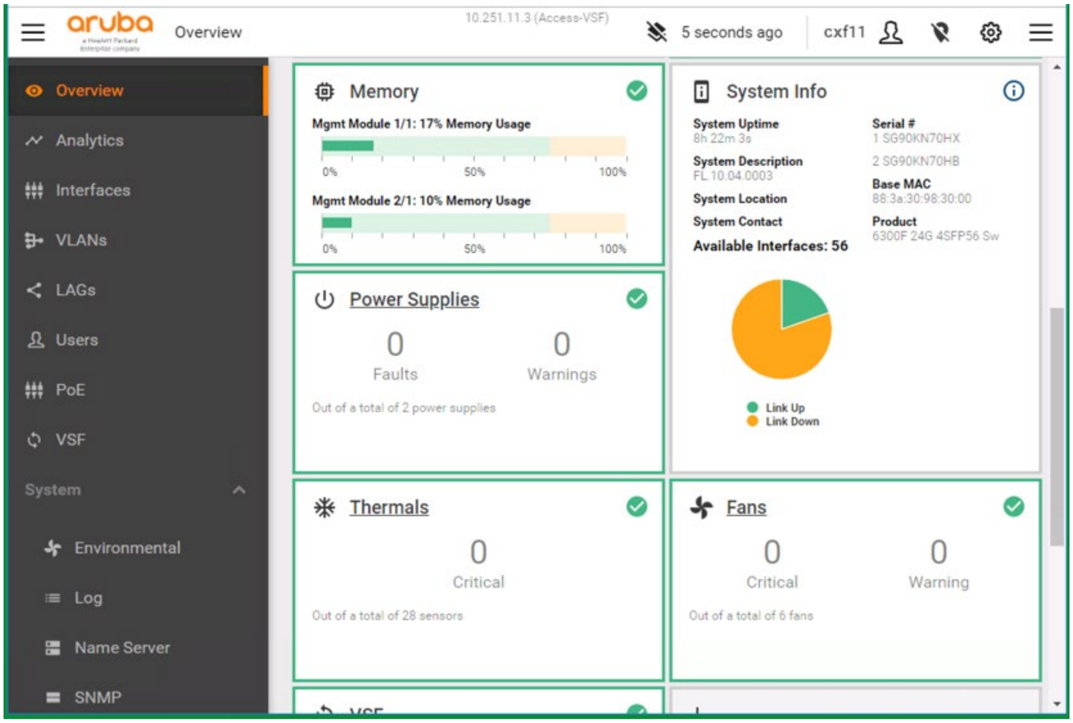

Mgmt Module 1/1: 17% Memory Usage

Mgmt Module 2/1: 10% Memory Usage

System Uptime
8h 22m 3s

System Description
FL.10.04.0003

System Location

System Contact

Available Interfaces: 56

Serial #
1 SG90KN70HX
2 SG90KN70HB

Base MAC
88:3a:30:98:30:00

Product
6300F 24G 4SFP56 Sw

Power Supplies: 0 Faults, 0 Warnings
Out of a total of 2 power supplies

Thermals: 0 Critical
Out of a total of 28 sensors

Fans: 0 Critical, 0 Warning
Out of a total of 6 fans

Figure 13-7: Overview 2

What is the Memory utilization on each stack member?

____

What are the serial number of both units?

____

What percentage of interfaces are down?

____

Is there any thermal or fan alarm?

____

6. Scroll down, then click the "**+**" sign in an open widget slot. It will ask for an interface number.
7. Select port **1/1/3**. In order to start monitoring the interface.

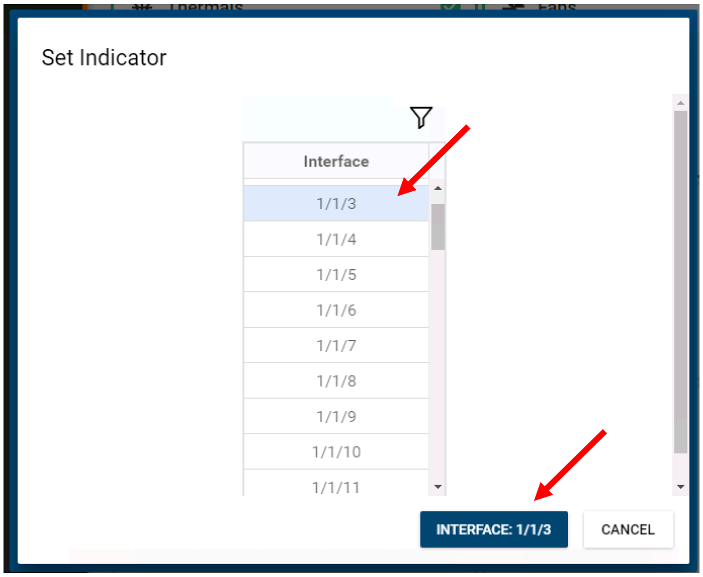

Figure 13-8: Physical Interfaces

8. Repeat step 7 with ports **1/1/21** and **2/1/21**, these are uplinks to Core-1.

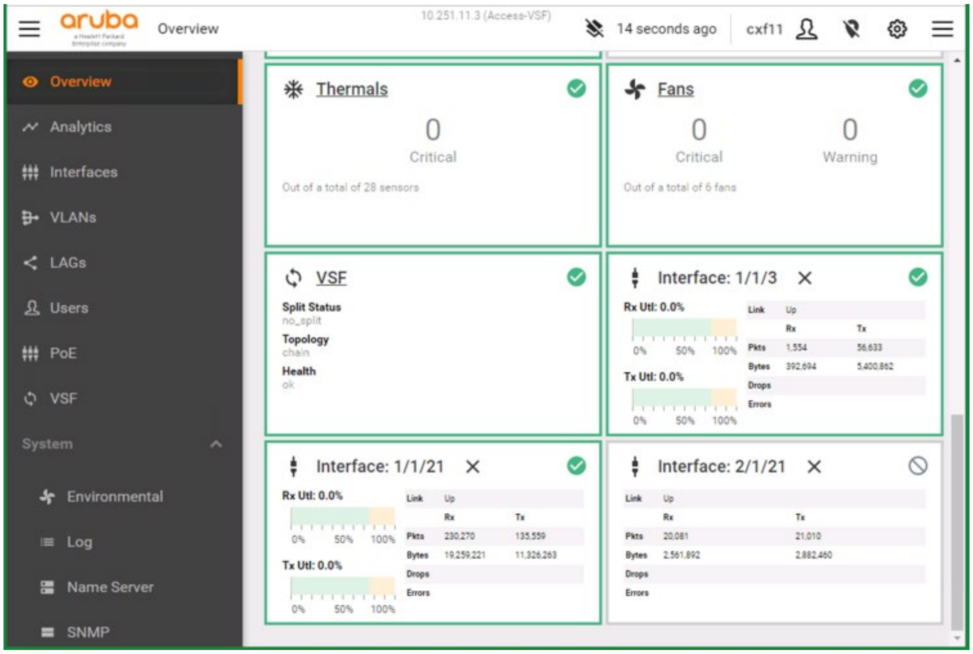

Thermals
0
Critical
Out of a total of 28 sensors

Fans
0 Critical
0 Warning
Out of a total of 6 fans

VSF
Split Status
no_split
Topology
chain
Health
ok

Interface: 1/1/3
Rx Utl: 0.0%
Tx Utl: 0.0%

Link	Up	
	Rx	Tx
Pkts	1,554	56,633
Bytes	392,694	5,400,862
Drops		
Errors		

Interface: 1/1/21
Rx Utl: 0.0%
Tx Utl: 0.0%

Link	Up	
	Rx	Tx
Pkts	230,270	135,559
Bytes	19,259,221	11,326,263
Drops		
Errors		

Interface: 2/1/21

Link	Up	
	Rx	Tx
Pkts	20,081	21,010
Bytes	2,561,892	2,882,460
Drops		
Errors		

Figure 13-9: Overview 3

What is the status of port 1/1/3?

____

Access-VSF: Member 1

9. Access Member 1’s console session.
10. Disable port **1/1/3**.

```
T11-Access-VSF# configure terminal
T11-Access-VSF(config)# interface 1/1/3
T11-Access-VSF(config-if)# shutdown
```

PC-1

11. Move back to the web session.

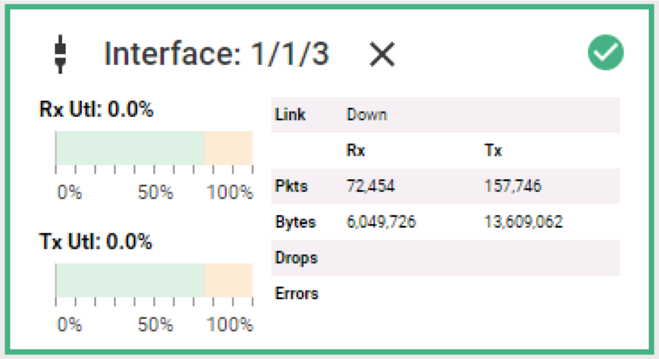

Interface: 1/1/3

Rx Utl: 0.0%

Tx Utl: 0.0%

	Rx	Tx
Link	Down	
Pkts	72,454	157,746
Bytes	6,049,726	13,609,062
Drops		
Errors		

Figure 13-10: Interface 1/1/3

Was there any change the link status?

Access-VSF: Member 1

12. Move back to Member 1.
13. Enable the port.

```
T11-Access-VSF(config-if)# no shutdown
```

PC-1

14. Move back to the web session.

Was is the VSF Split status?

Was is the VSF topology?

Was is the VSF Health?

---

15. Click at the **VSF** hyper-link. That will take you to the VSF page.

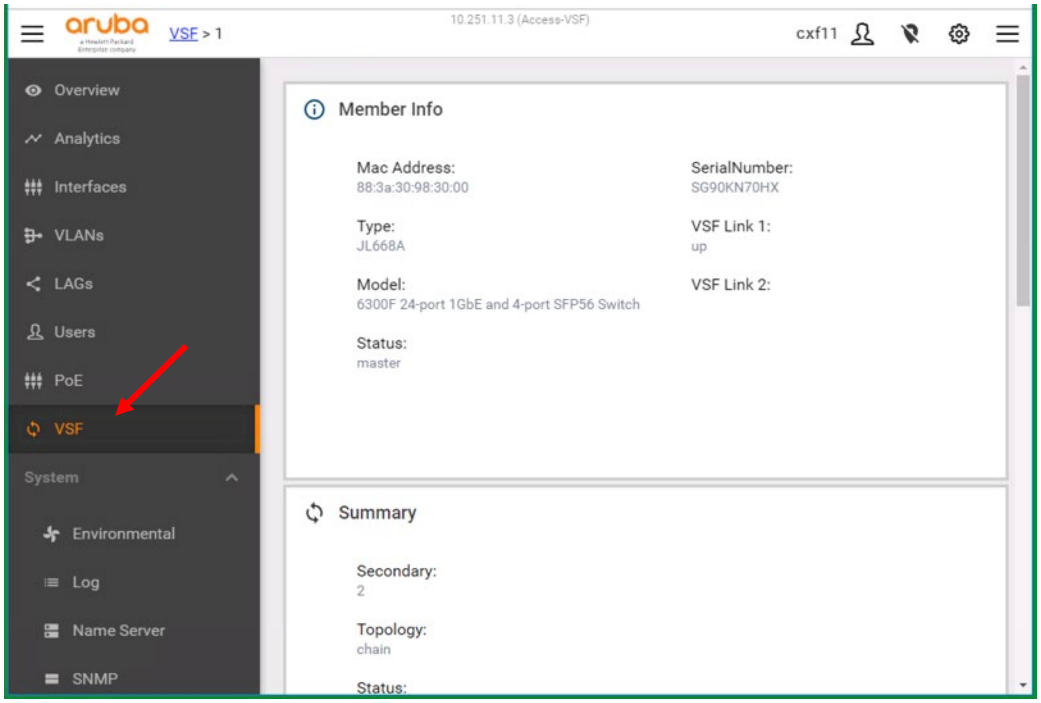

Member Info

Mac Address:
88:3a:30:98:30:00

SerialNumber:
SG90KN70HX

Type:
JL668A

VSF Link 1:
up

Model:
6300F 24-port 1GbE and 4-port SFP56 Switch

VSF Link 2:

Status:
master

Summary

Secondary:
2

Topology:
chain

Status:

Figure 13-11: VSF

Whose information is shown member 1 or 2?

---

16. Scroll down.

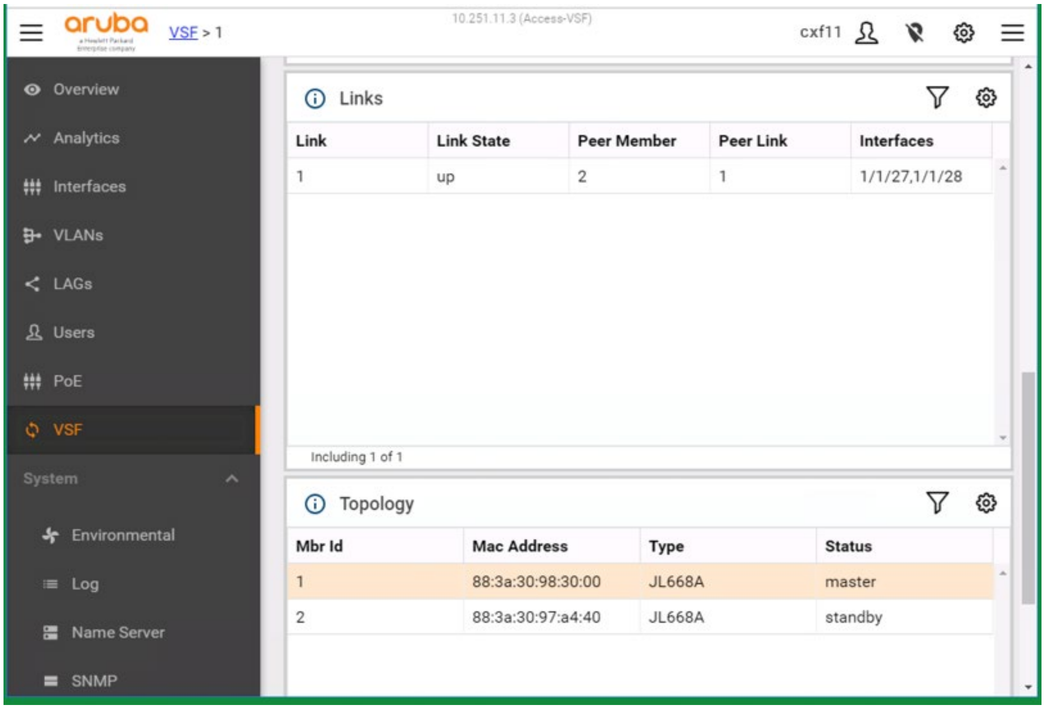

Links

Link	Link State	Peer Member	Peer Link	Interfaces
1	up	2	1	1/1/27,1/1/28

Including 1 of 1

Topology

Mbr Id	Mac Address	Type	Status
1	88:3a:30:98:30:00	JL668A	master
2	88:3a:30:97:a4:40	JL668A	standby

Figure 13-12: VSF 2

What physical ports are being used for the logical VSF link?

____

17. Select member 2 in the topology table.

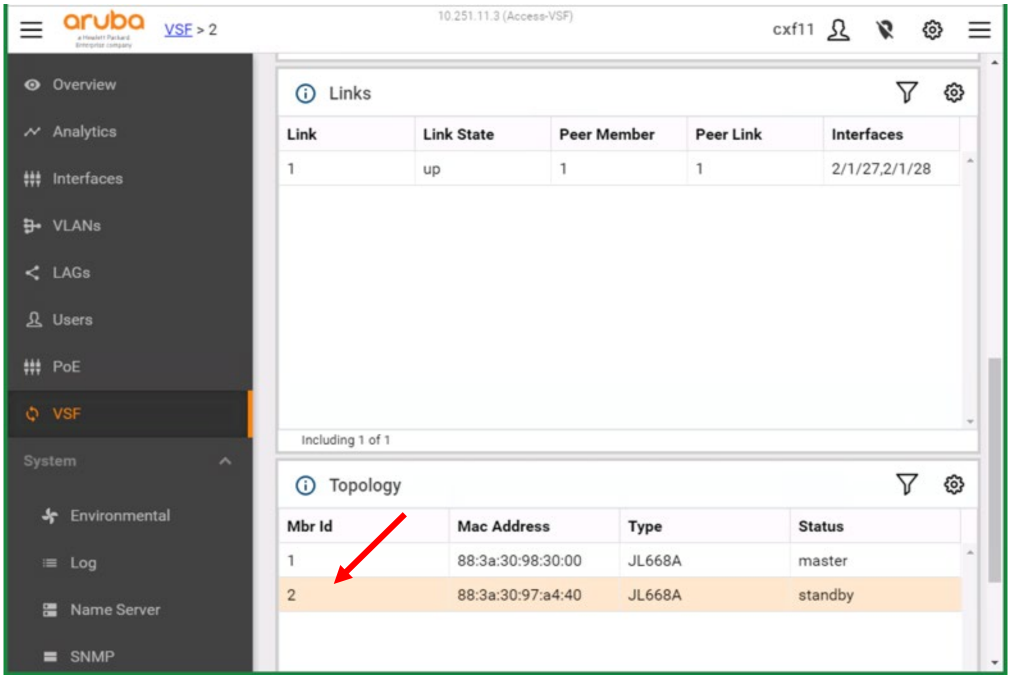

Links

Link	Link State	Peer Member	Peer Link	Interfaces
1	up	1	1	2/1/27,2/1/28

Including 1 of 1

Topology

Mbr Id	Mac Address	Type	Status
1	88:3a:30:98:30:00	JL668A	master
2	88:3a:30:97:a4:40	JL668A	standby

Figure 13-13: VSF 3

What physical ports are being used for the logical VSF link?

____

18. Click on **Interfaces** in the left navigation pane.

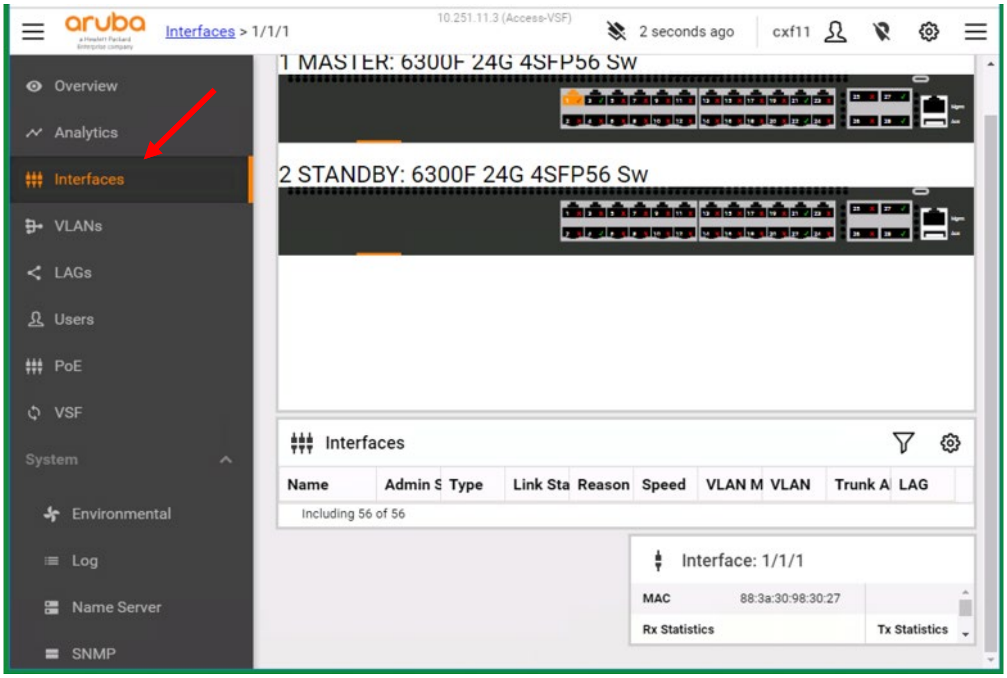

Figure 13-14: Front pane

How many ports do the switches have?

____

What interfaces are up on member 1?

____

What interfaces are up on member 2?

____

19. Click on **VLANs** in the left navigation pane.

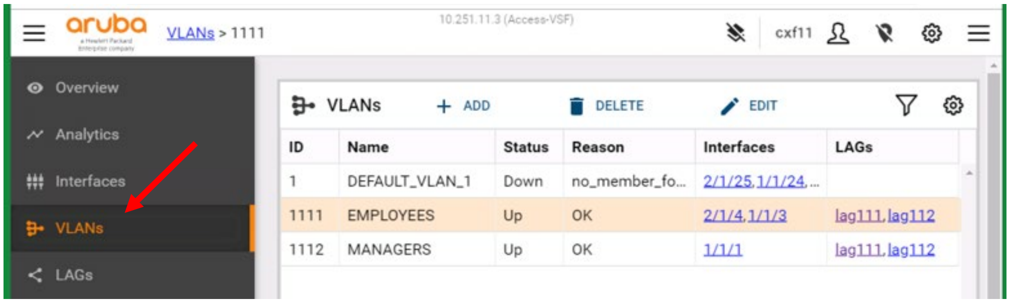

ID	Name	Status	Reason	Interfaces	LAGs
1	DEFAULT_VLAN_1	Down	no_member_fo...	2/1/25,1/1/24,...	
1111	EMPLOYEES	Up	OK	2/1/4,1/1/3	lag111,lag112
1112	MANAGERS	Up	OK	1/1/1	lag111,lag112

Figure 13-15: VLANs

How many VLANs are listed and what are their names?

What ports are members of VLAN X12?

---

20. Click on **LAGs** in the left navigation pane.

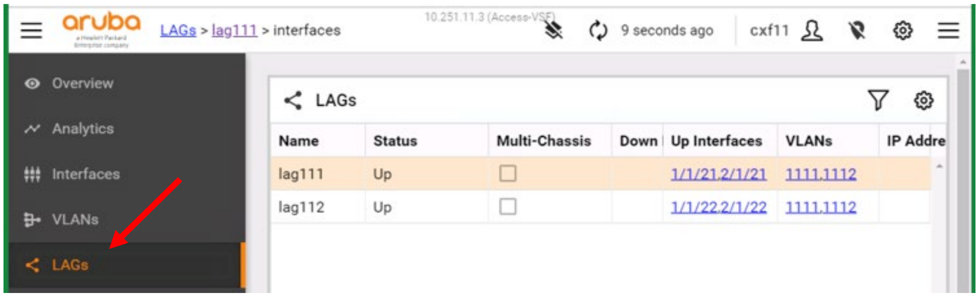

Name	Status	Multi-Chassis	Down	Up Interfaces	VLANs	IP Addre
lag111	Up	□		1/1/21,2/1/21	1111,1112	
lag112	Up	□		1/1/22,2/1/22	1111,1112	

Figure 13-16: LAGs

How many LAGs are created?

Ports are used in these LAGs?

---

21. Expand "System" on the left navigation pane. Then click on "**Environmental**".

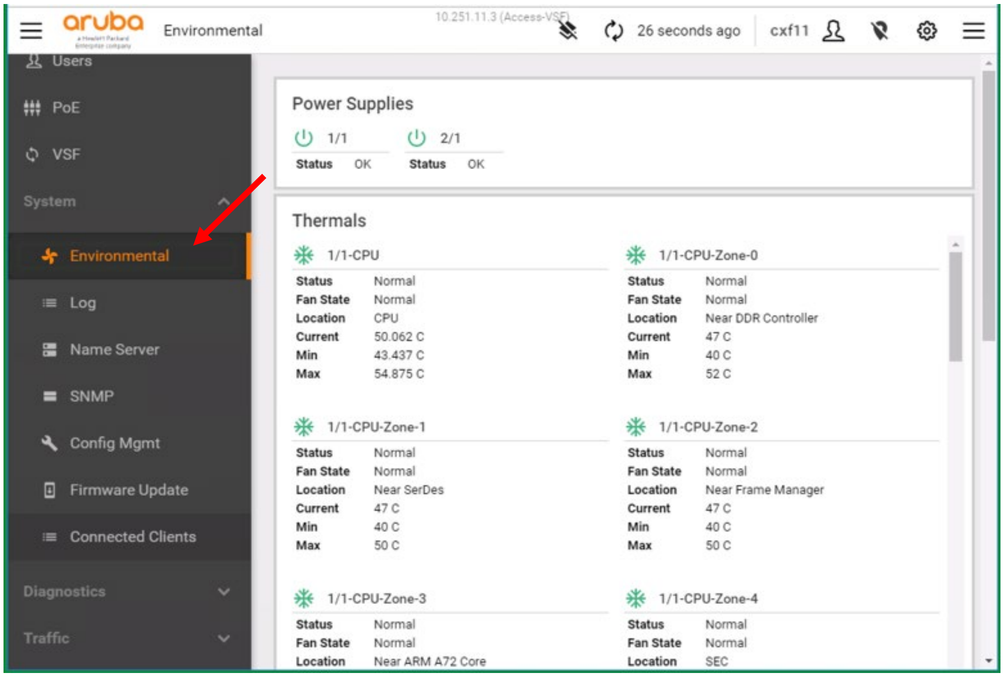

Power Supplies

1/1
Status OK

2/1
Status OK

Thermals

1/1-CPU
Status Normal
Fan State Normal
Location CPU
Current 50.062 C
Min 43.437 C
Max 54.875 C

1/1-CPU-Zone-0
Status Normal
Fan State Normal
Location Near DDR Controller
Current 47 C
Min 40 C
Max 52 C

1/1-CPU-Zone-1
Status Normal
Fan State Normal
Location Near SerDes
Current 47 C
Min 40 C
Max 50 C

1/1-CPU-Zone-2
Status Normal
Fan State Normal
Location Near Frame Manager
Current 47 C
Min 40 C
Max 50 C

1/1-CPU-Zone-3
Status Normal
Fan State Normal
Location Near ARM A72 Core

1/1-CPU-Zone-4
Status Normal
Fan State Normal
Location SEC

Figure 13-17: System - Environmental

How many power supplies does the stack have?

____

What is the current temperature of CPU 1/1?

____

What is the current temperature of CPU 1/1 Zones 0 to 4?

____

____

____

22. Click at **System -> Log**.

23. Select any of the entries.

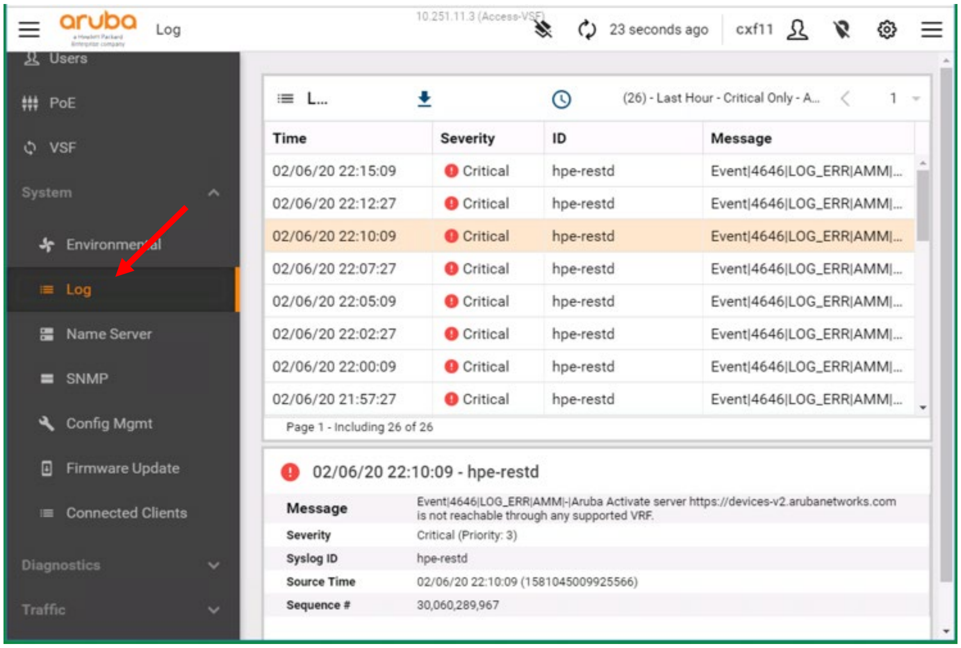

Time	Severity	ID	Message
02/06/20 22:15:09	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:12:27	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:10:09	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:07:27	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:05:09	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:02:27	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 22:00:09	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...
02/06/20 21:57:27	Critical	hpe-restd	Event|4646|LOG_ERR|AMM|...

Page 1 - Including 26 of 26

02/06/20 22:10:09 - hpe-restd

Message	Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-v2.arubanetworks.com is not reachable through any supported VRF.
Severity	Critical (Priority: 3)
Syslog ID	hpe-restd
Source Time	02/06/20 22:10:09 (1581045009925566)
Sequence #	30,060,289,967

Figure 13-18: System - Log

What is the severity of the log record?

____

What is the message of the log record?

____

____

NOTE: It is likely the log refers to the lack of connectivity to Aruba Activate. The switches do not currently have internet access.

NOTE: Aruba Activate provides Zero Touch Provisioning and can facilitate centralized management platforms.

24. Click on **System -> Connected Clients**, then scroll down. This shows the LLDP table with all discovered neighbors.

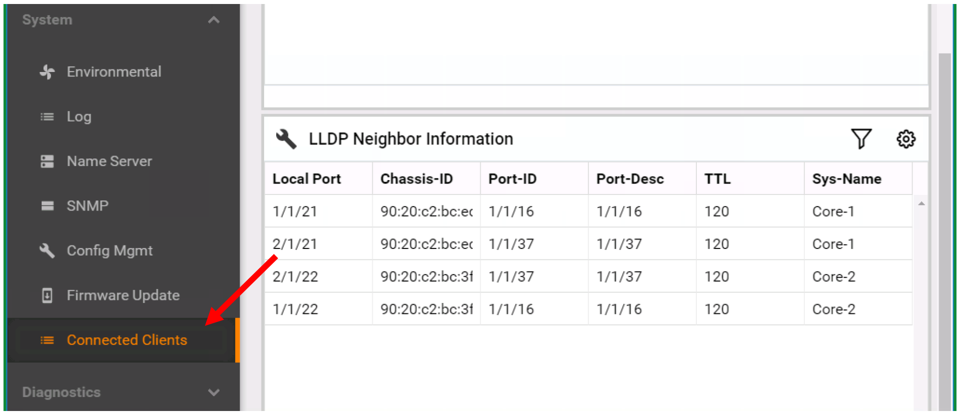

LLDP Neighbor Information

Local Port	Chassis-ID	Port-ID	Port-Desc	TTL	Sys-Name
1/1/21	90:20:c2:bc:ec	1/1/16	1/1/16	120	Core-1
2/1/21	90:20:c2:bc:ec	1/1/37	1/1/37	120	Core-1
2/1/22	90:20:c2:bc:3f	1/1/37	1/1/37	120	Core-2
1/1/22	90:20:c2:bc:3f	1/1/16	1/1/16	120	Core-2

Figure 13-19: System - Connected Clients

25. Expand Diagnostics, then click on Ping.

26. Type **10.251.X.200** as IPv4 Target, then check **Use Management Interface** checkbox". This IP address is owned by the NETEDIT system you will use in next lab.

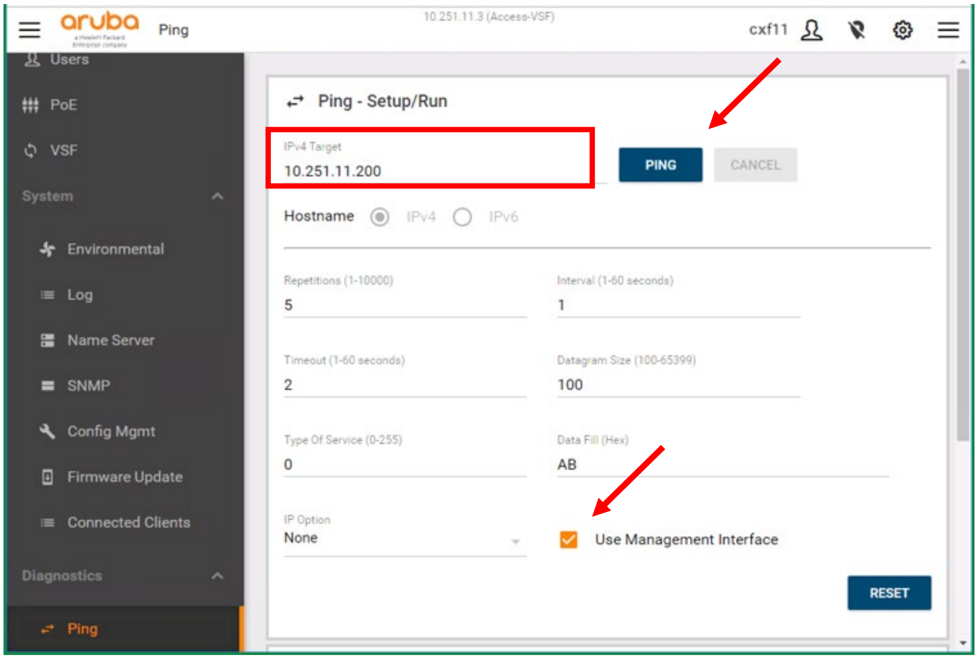

Ping - Setup/Run

IPv4 Target: 10.251.11.200

Hostname ● IPv4 ○ IPv6

Repetitions (1-10000): 5

Interval (1-60 seconds): 1

Timeout (1-60 seconds): 2

Datagram Size (100-65399): 100

Type Of Service (0-255): 0

Data Fill (Hex): AB

IP Option: None

☑ Use Management Interface

Figure 13-20: Diagnostics - Ping

27. Press the **PING** button and wait. Be patient…

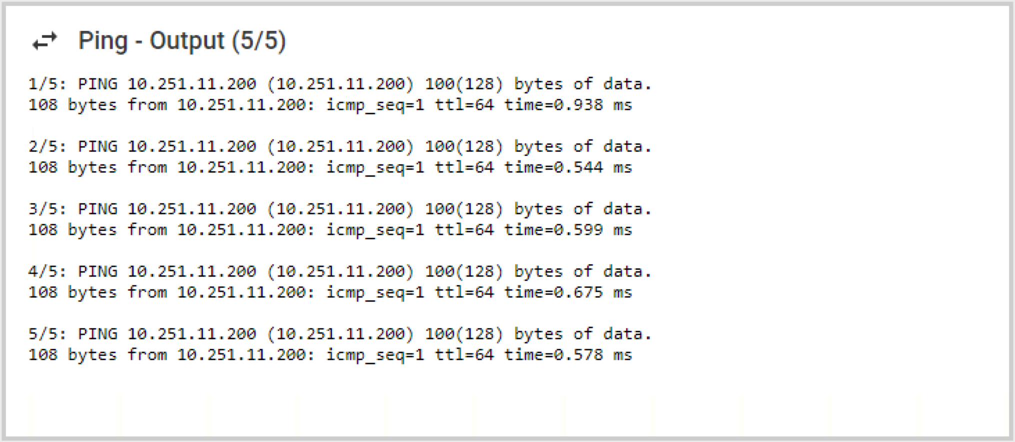

```
1/5: PING 10.251.11.200 (10.251.11.200) 100(128) bytes of data.
108 bytes from 10.251.11.200: icmp_seq=1 ttl=64 time=0.938 ms

2/5: PING 10.251.11.200 (10.251.11.200) 100(128) bytes of data.
108 bytes from 10.251.11.200: icmp_seq=1 ttl=64 time=0.544 ms

3/5: PING 10.251.11.200 (10.251.11.200) 100(128) bytes of data.
108 bytes from 10.251.11.200: icmp_seq=1 ttl=64 time=0.599 ms

4/5: PING 10.251.11.200 (10.251.11.200) 100(128) bytes of data.
108 bytes from 10.251.11.200: icmp_seq=1 ttl=64 time=0.675 ms

5/5: PING 10.251.11.200 (10.251.11.200) 100(128) bytes of data.
108 bytes from 10.251.11.200: icmp_seq=1 ttl=64 time=0.578 ms
```

Figure 13-21: Ping result

Was the ping successful?

____

28. Go to **Diagnostics -> Show Tech**.
29. Click on **GENERATE**. This will create the “Show Tech” support file.
30. Click on **EXPORT**. This will download the file through the browser. File will show up at the bottom of the browser.

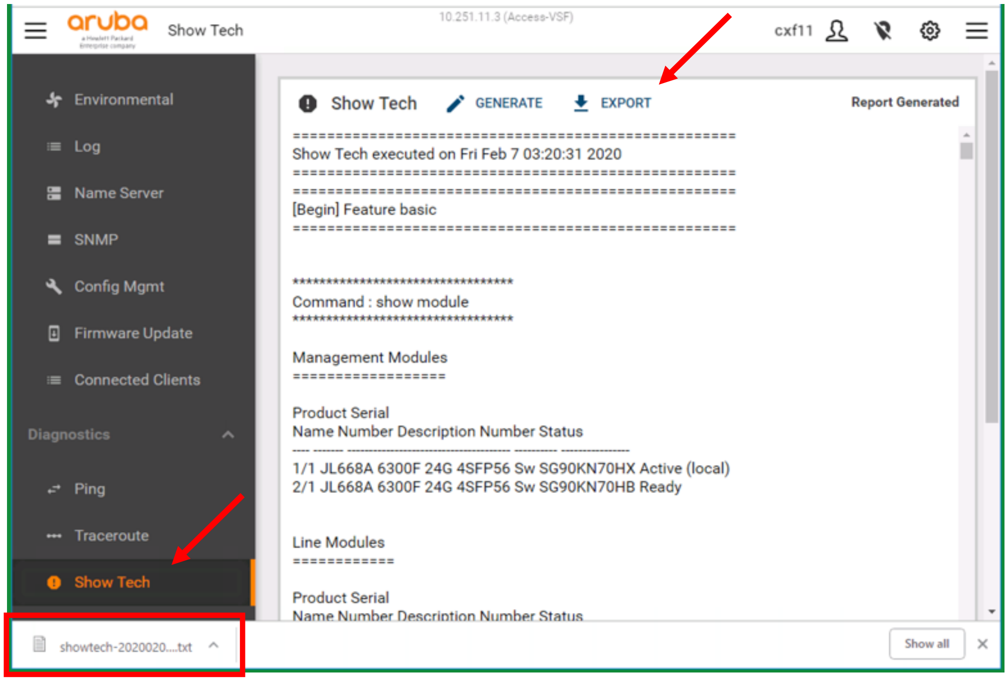

```
====================================================
Show Tech executed on Fri Feb 7 03:20:31 2020
====================================================
====================================================
[Begin] Feature basic
====================================================


**********************************
Command : show module
**********************************

Management Modules
==================

Product Serial
Name Number Description Number Status
---- ------- ---------------------------------- ---------- ----------------
1/1 JL668A 6300F 24G 4SFP56 Sw SG90KN70HX Active (local)
2/1 JL668A 6300F 24G 4SFP56 Sw SG90KN70HB Ready


Line Modules
============

Product Serial
Name Number Description Number Status
```

Figure 13-22: Diagnostics – Show Tech

NOTE: When opening a Technical Assistance Center (TAC) support case, one of the pieces of information they will first ask for is this output. It is always a good practice to generate it and download it in advance.

31. Click on the gear icon in the top right corner, then select **V10.04 API**. This will open another browser tab and display the AOS-CX REST API documentation.

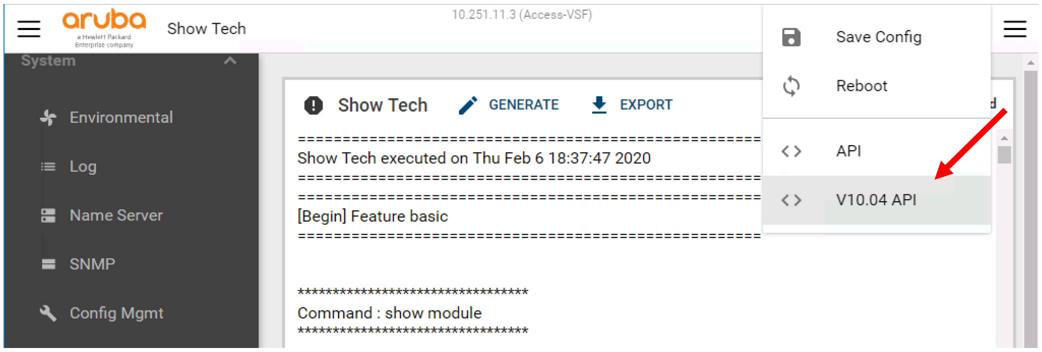

```
==========================================================
Show Tech executed on Thu Feb 6 18:37:47 2020
==========================================================
==========================================================
[Begin] Feature basic
==========================================================

*********************************
Command : show module
*********************************
```

Figure 13-23: Menu - API

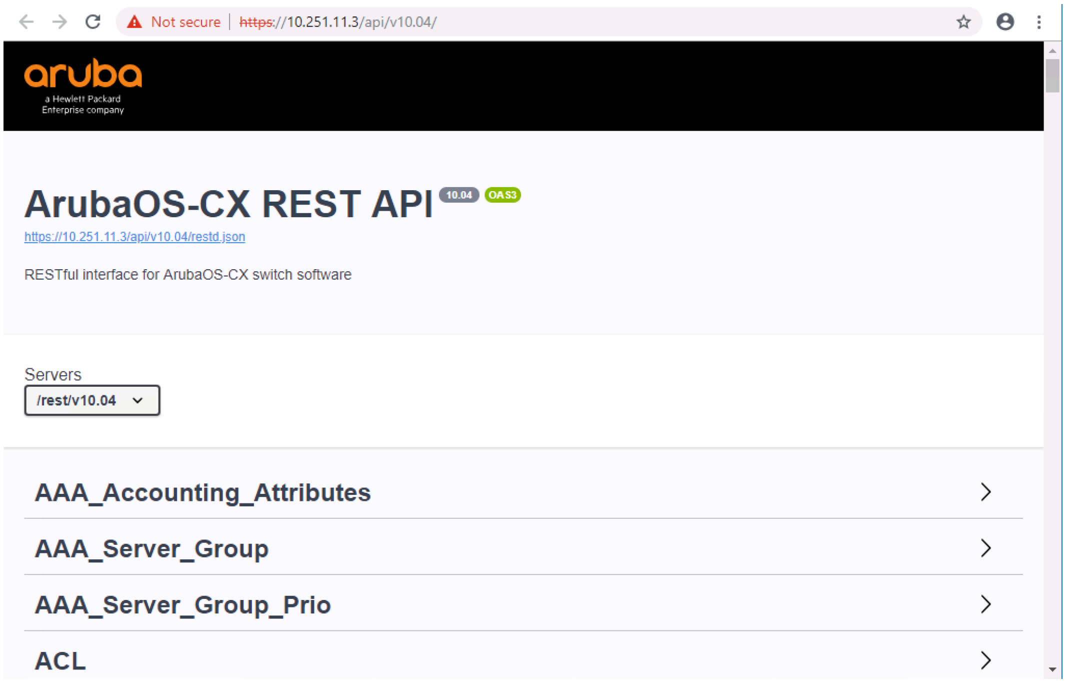
Not secure
https://10.251.11.3/api/v10.04/
aruba a Hewlett Packard Enterprise company

ArubaOS-CX REST API 10.04 OAS3

https://10.251.11.3/api/v10.04/restd.json

RESTful interface for ArubaOS-CX switch software

Servers
/rest/v10.04

AAA_Accounting_Attributes

AAA_Server_Group

AAA_Server_Group_Prio

ACL

Figure 13-24: AOS-CX REST API

NOTE: Switches running the AOS-CX software are fully programmable with a REST (Representational State Transfer) API, allowing easy integration with other devices both on premises and in the cloud. This programmability, combined with the Aruba Network Analytics Engine, accelerates a network administrator's understanding and response to network issues.

The AOS-CX REST API enables programmatic access to the AOS-CX configuration and state database at the heart of the switch. By using a structured model, changes to the content and formatting of the CLI output do not affect the programs you write. And, because the configuration is stored in a structured database instead of a text file, rolling back changes is easy. This reduces the risk of downtime and performance issues.

You will now access the Web UI of Core-1 and see the minor differences between a 8325 switch and a 6300 switch.

32. Open another browser tab.
33. In the URL field type the management IP address of Core-1 (**10.251.X.201**).

What Navigation Pane option is different to the UI of the 6300 switch?

____

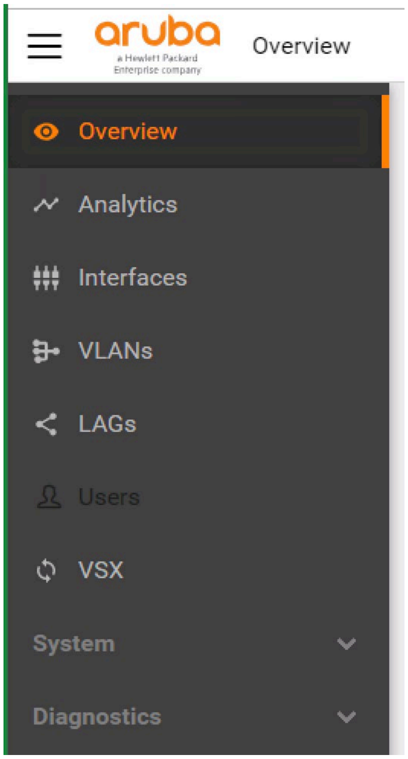

Figure 13-25: Navigation Pane

34. Click on **Interfaces**.

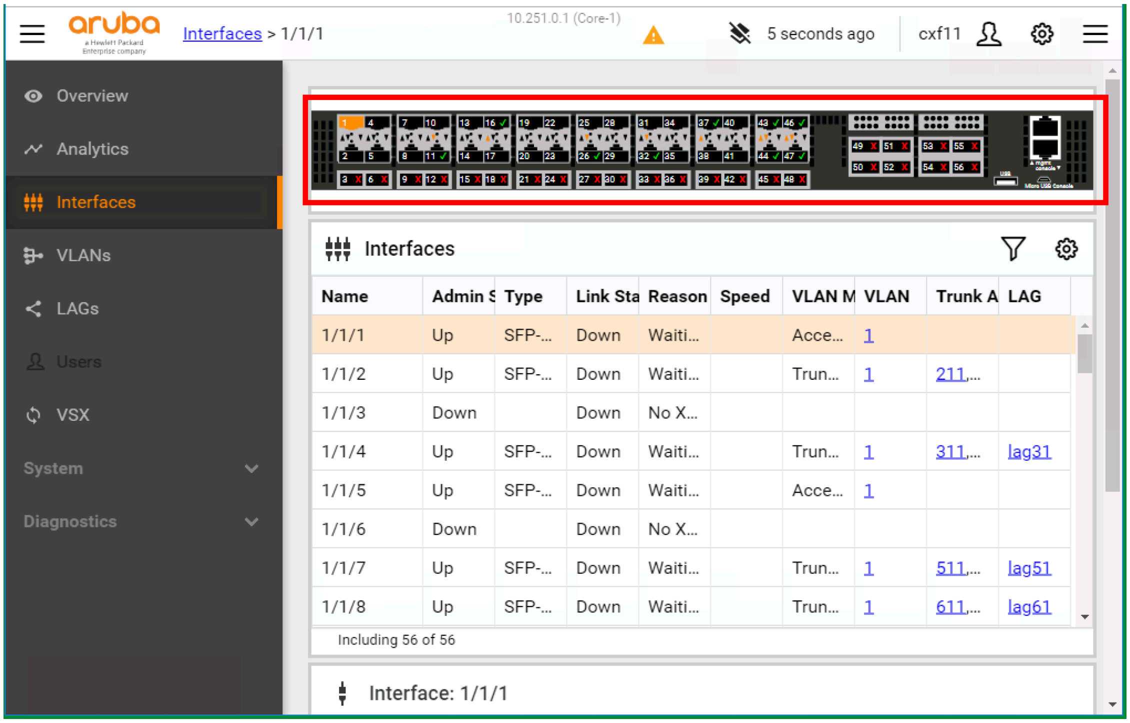

Interfaces

Name	Admin S	Type	Link Sta	Reason	Speed	VLAN M	VLAN	Trunk A	LAG
1/1/1	Up	SFP-...	Down	Waiti...		Acce...	1		
1/1/2	Up	SFP-...	Down	Waiti...		Trun...	1	211,...	
1/1/3	Down		Down	No X...					
1/1/4	Up	SFP-...	Down	Waiti...		Trun...	1	311,...	lag31
1/1/5	Up	SFP-...	Down	Waiti...		Acce...	1		
1/1/6	Down		Down	No X...					
1/1/7	Up	SFP-...	Down	Waiti...		Trun...	1	511,...	lag51
1/1/8	Up	SFP-...	Down	Waiti...		Trun...	1	611,...	lag61

Including 56 of 56

Interface: 1/1/1

Figure 13-26: 8325 Interfaces

What differences can you see to the panel shown in the 6300’s UI (step 18)?

____

Task 5: Save Your Configurations

Objectives

You will now proceed to create a checkpoint, save your configuration and download it as a file in order to keep a backup of the current configuration.

Steps

PC-1

1. Move back to the browser tab of the 6300’s UI, you might need to login using “**cxfX/aruba123**”
2. Click on **Config Mgmt**.

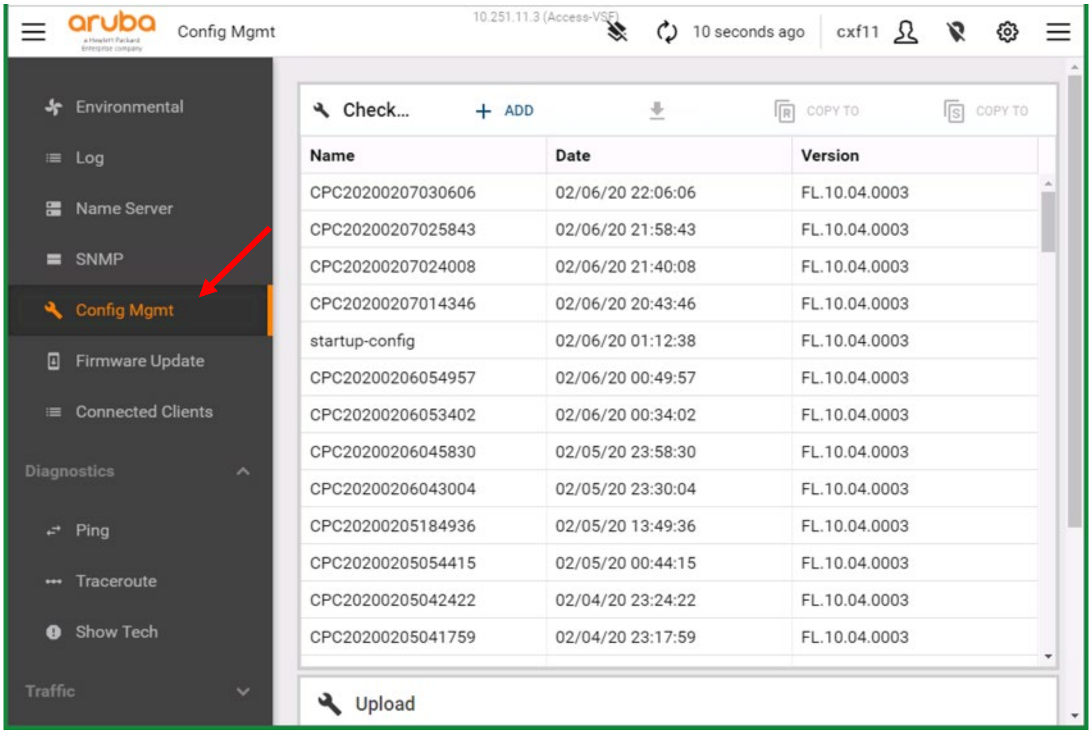

Name	Date	Version
CPC20200207030606	02/06/20 22:06:06	FL.10.04.0003
CPC20200207025843	02/06/20 21:58:43	FL.10.04.0003
CPC20200207024008	02/06/20 21:40:08	FL.10.04.0003
CPC20200207014346	02/06/20 20:43:46	FL.10.04.0003
startup-config	02/06/20 01:12:38	FL.10.04.0003
CPC20200206054957	02/06/20 00:49:57	FL.10.04.0003
CPC20200206053402	02/06/20 00:34:02	FL.10.04.0003
CPC20200206045830	02/05/20 23:58:30	FL.10.04.0003
CPC20200206043004	02/05/20 23:30:04	FL.10.04.0003
CPC20200205184936	02/05/20 13:49:36	FL.10.04.0003
CPC20200205054415	02/05/20 00:44:15	FL.10.04.0003
CPC20200205042422	02/04/20 23:24:22	FL.10.04.0003
CPC20200205041759	02/04/20 23:17:59	FL.10.04.0003

Figure 13-27: Config Mgmt

3. Click on **ADD**.

4. Type **Lab13_final** as the checkpoint name, then click on **Create Checkpoint**, and **close** when done.

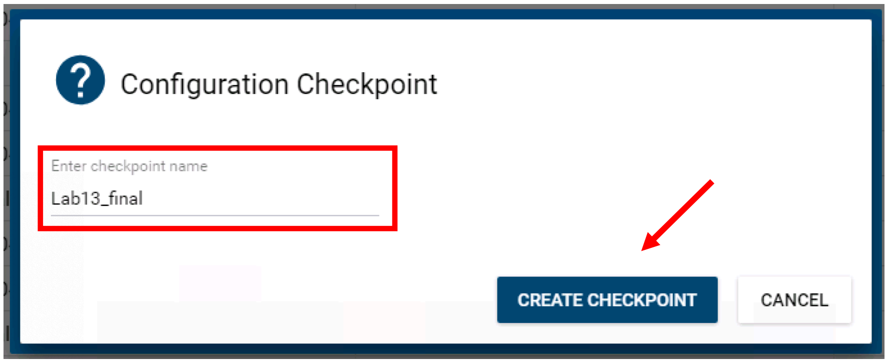

Configuration Checkpoint

Enter checkpoint name: Lab13_final

Figure 13-28: Configuration checkpoint

5. Select Lab13_final checkpoint, then click on **Copy to Startup** button.

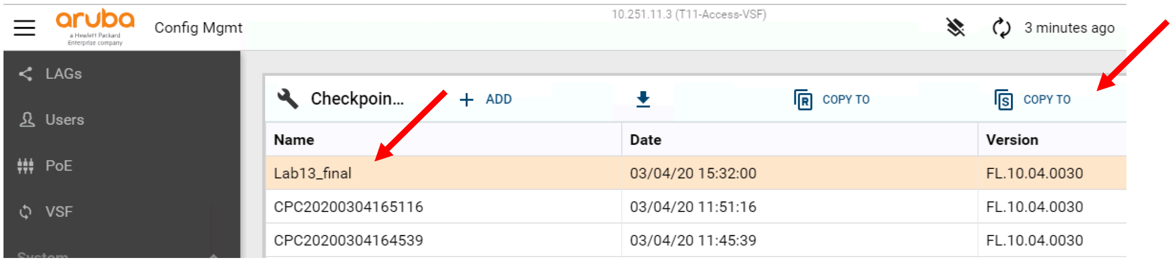

Name	Date	Version
Lab13_final	03/04/20 15:32:00	FL.10.04.0030
CPC20200304165116	03/04/20 11:51:16	FL.10.04.0030
CPC20200304164539	03/04/20 11:45:39	FL.10.04.0030

Figure 13-29: Configuration checkpoint

6. Click on **Copy** button then close.

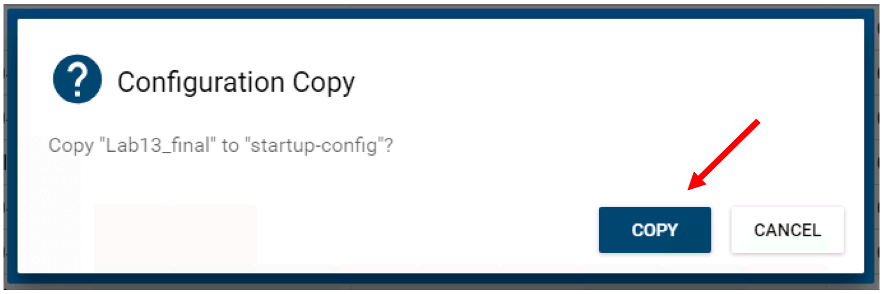

Configuration Copy

Copy "Lab13_final" to "startup-config"?

Figure 13-30: Configuration Copy

7. Select **Lab13_final** checkpoint then click on the **download** button, the backup will show up at the bottom of the browser.

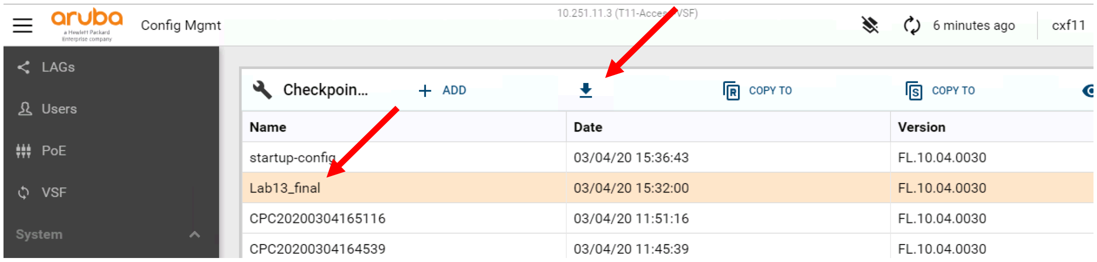

Name	Date	Version
startup-config	03/04/20 15:36:43	FL.10.04.0030
Lab13_final	03/04/20 15:32:00	FL.10.04.0030
CPC20200304165116	03/04/20 11:51:16	FL.10.04.0030
CPC20200304164539	03/04/20 11:45:39	FL.10.04.0030

8. Click on **Close** button.

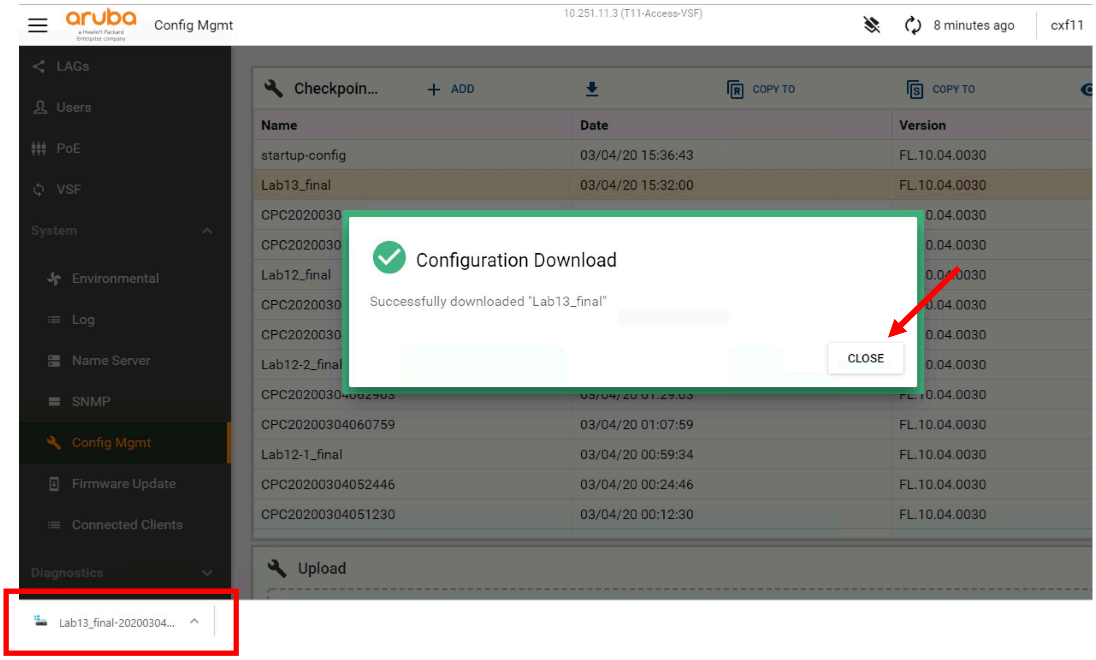

Name	Date	Version
startup-config	03/04/20 15:36:43	FL.10.04.0030
Lab13_final	03/04/20 15:32:00	FL.10.04.0030
CPC2020030		0.04.0030
CPC2020030		0.04.0030
Lab12_final		0.04.0030
CPC2020030		0.04.0030
CPC2020030		0.04.0030
Lab12-2_final		0.04.0030
CPC20200304062903	03/04/20 01:29:03	FL.10.04.0030
CPC20200304060759	03/04/20 01:07:59	FL.10.04.0030
Lab12-1_final	03/04/20 00:59:34	FL.10.04.0030
CPC20200304052446	03/04/20 00:24:46	FL.10.04.0030
CPC20200304051230	03/04/20 00:12:30	FL.10.04.0030

Configuration Download

Successfully downloaded "Lab13_final"

CLOSE

Figure 13-31: Configuration Download

9. Click on the **Close** button.

NOTICE: This is the last configuration step you will apply on this course lab. Feel free to ask your instructor for this backup, he will collect it and share it with you.

You have completed Lab 13!

ArubaOS-CX Switching Fundamentals

Lab 14: Monitoring Devices with Aruba NetEdit

Overview

After enabling remote management, BigStartup's IT staff wants to know if there is a way to monitor and manage both the main office and the remote locations from a single pane of glass. Currently they are opening individual web sessions to each switch. To provide a consolidated monitoring and management service, you will demonstrate Aruba NetEdit.

Objectives

After completing this lab, you will be able to:

- Access NetEdit and update admin credentials
- Discover switching devices
- Monitor switching devices
- Run a Deployment Plan
- Commit a deployment
- Inspect the logs

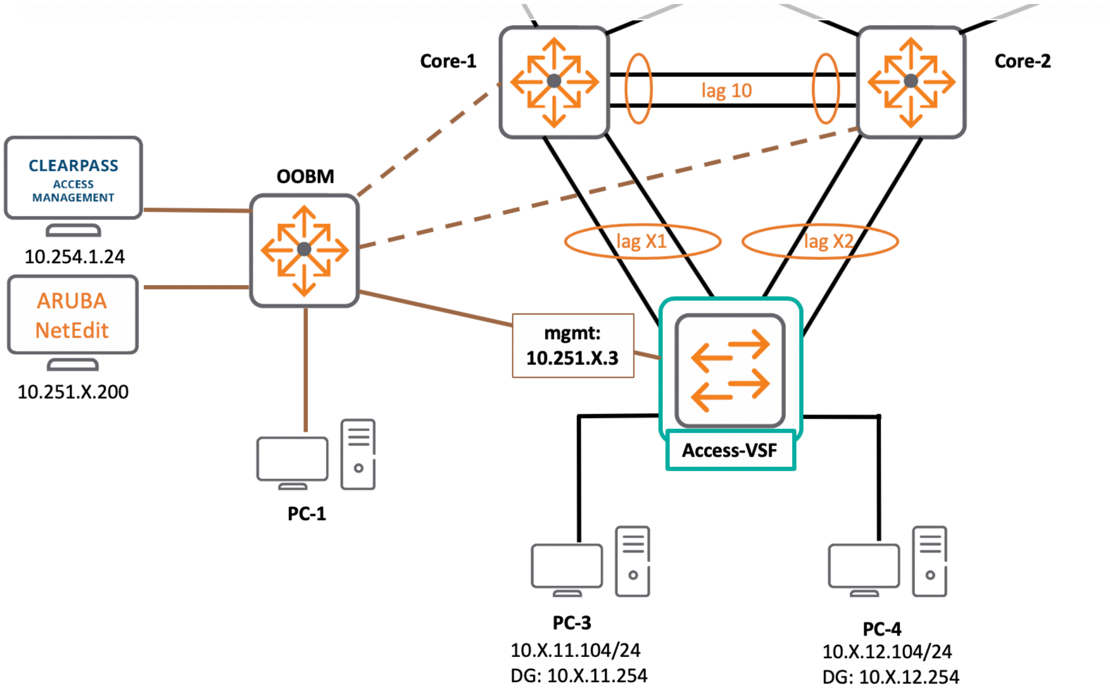

Figure 14-1: Lab Topology

Task 1: Discovering Devices in NetEdit

Objectives

In this lab you will access NetEdit for the first time, therefore you will be asked to update the “admin” credentials. Once inside you are going to add devices into its management database and proceed with regular monitoring and exploration tasks.

Steps

PC-1

1. Access PC-1.
2. Open a browser and type the NetEdit IP address in the URL field (**10.251.X.200**) then hit **[Enter]**.
3. Login with **admin** and no password (leave the field empty). You will be asked to change your password.

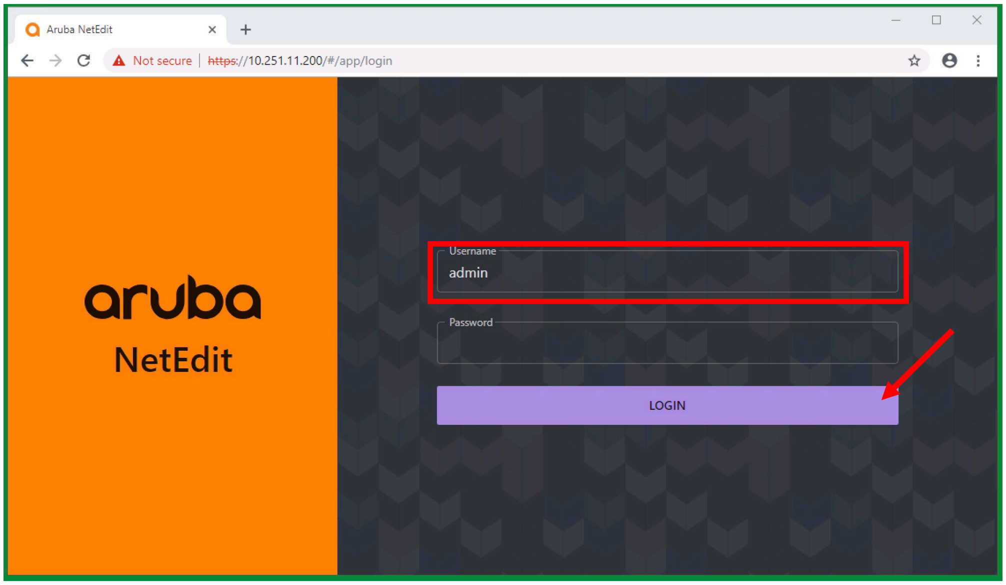

Figure 14-2: NetEdit login page

4. In the Password Change Required dialog box type “**aruba123**” with no quotes under the Password and Confirm Password fields.

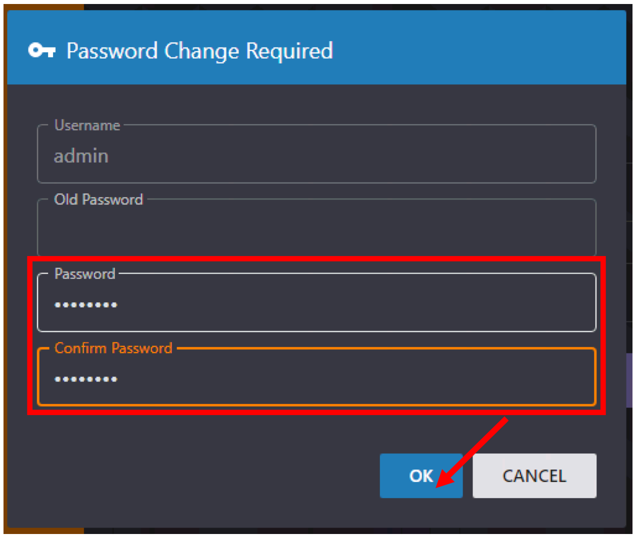

Figure 14-3: Update admin credentials

5. Click the **OK** button. That will take you to NetEdit **Overview**.

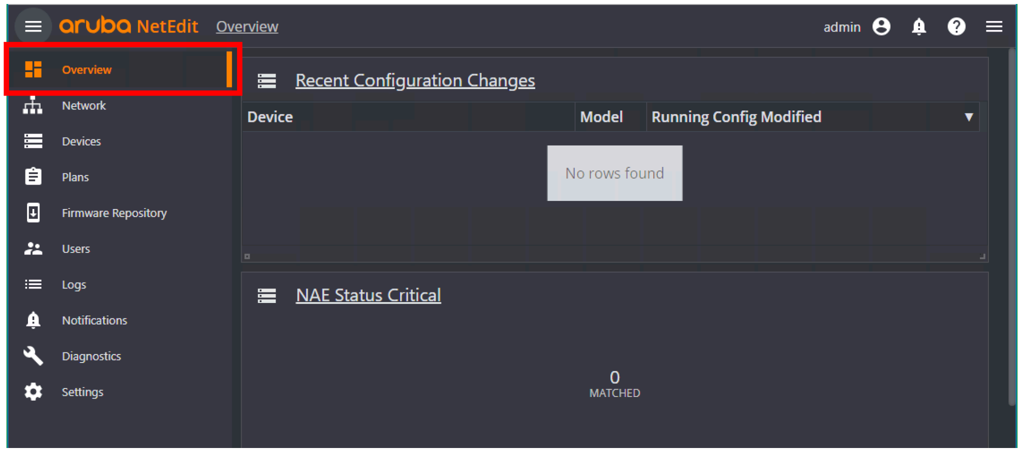

Figure 14-4: NetEdit Overview

6. In the left navigation panel click on **Devices**.
7. At the far right click on the **Action** button, then select **Discover Devices** from the menu that appears. A dialog box will show up.

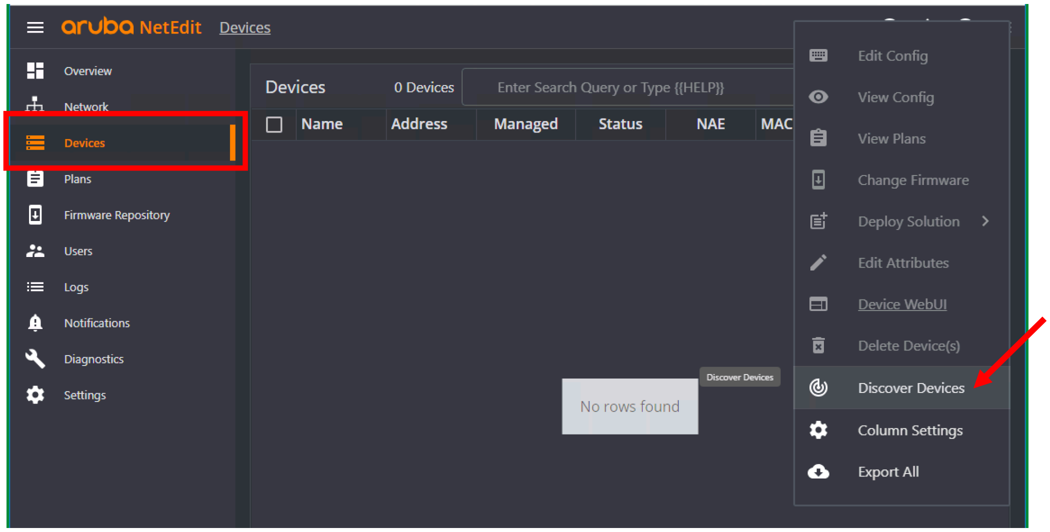

Figure 14-5: NetEdit Devices

8. Under Subnet type **10.251.X.3/32**.
9. Click on **Add Credentials**. A new dialog box appears.

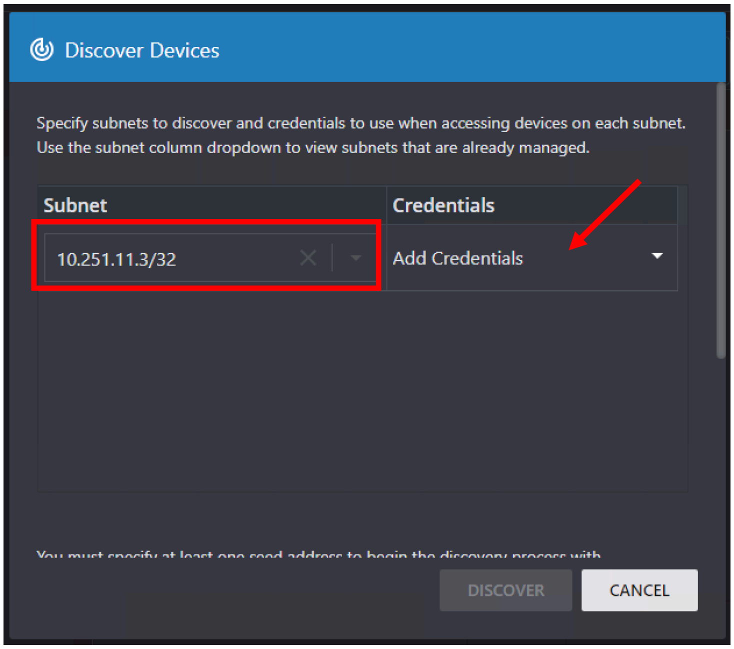

Specify subnets to discover and credentials to use when accessing devices on each subnet. Use the subnet column dropdown to view subnets that are already managed.

Subnet	Credentials
10.251.11.3/32	Add Credentials

You must specify at least one seed address to begin the discovery process with

Figure 14-6: Discover Devices

10. On credentials Name type **cxfX**.
11. Expand “REST – required for AOS-CX devices”, then type **cxfX** as username and **aruba123** as password.

12. Repeat **step 11** under **SSH – required for Change Validation**.
13. Click on the **eye** icon to confirm the password.

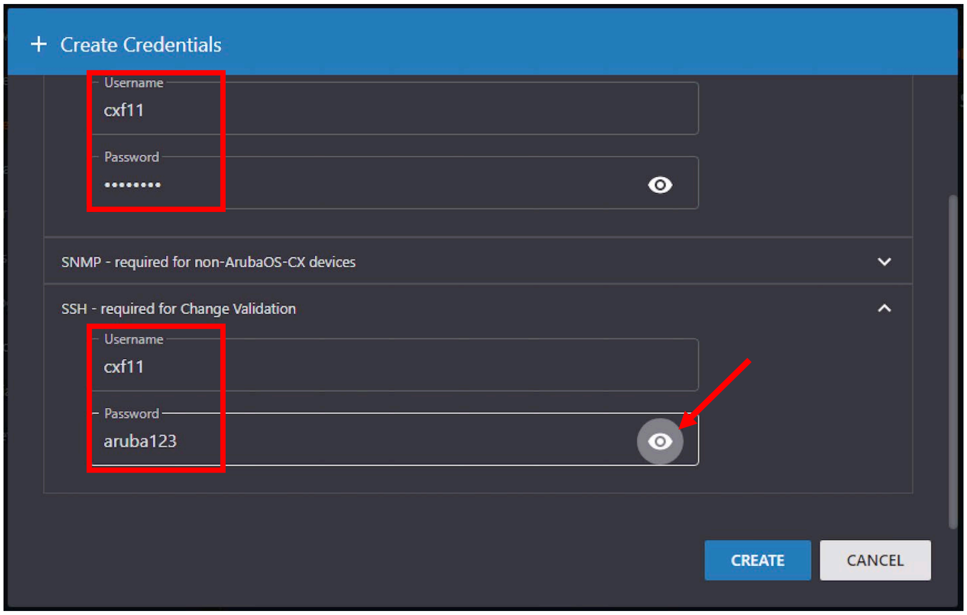

+ Create Credentials

Username: cxf11

Password: ••••••••

SNMP - required for non-ArubaOS-CX devices

SSH - required for Change Validation

Username: cxf11

Password: aruba123

Figure 14-7: Create Credentials

14. Click **CREATE** button.
15. Back in the Discover Devices dialog box, scroll down, then in the Seed Addresses area click the "**+**" sign. A new dialog box will show up.
16. Type **10.251.X.3** then click **ADD** button.

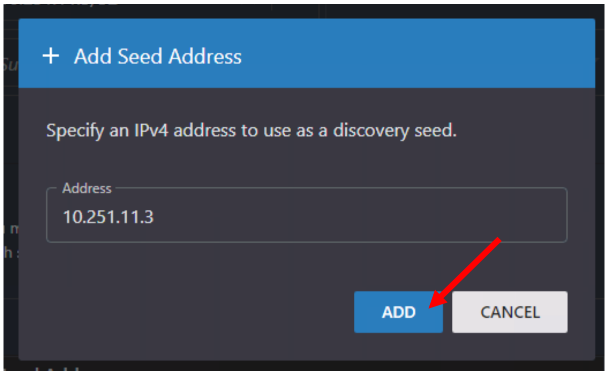

+ Add Seed Address

Specify an IPv4 address to use as a discovery seed.

Address: 10.251.11.3

Figure 14-8: Add Seed Address

17. Check the newly added Seed Address then click the **Discover** button.

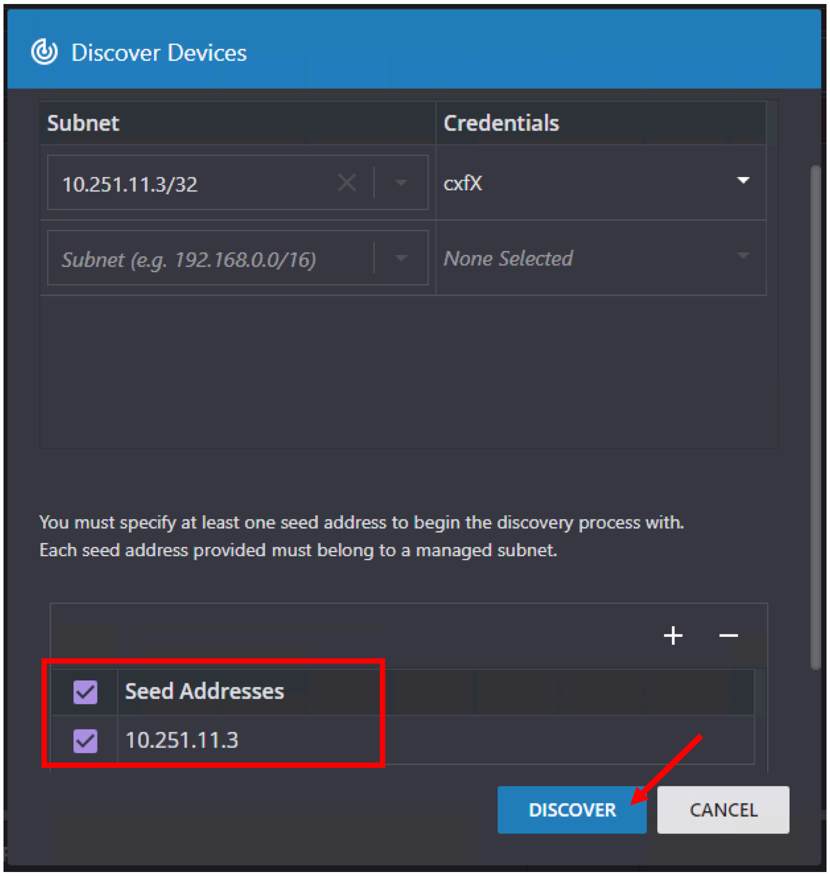

Subnet	Credentials
10.251.11.3/32	cxfX
	None Selected

You must specify at least one seed address to begin the discovery process with. Each seed address provided must belong to a managed subnet.

☑	Seed Addresses
☑	10.251.11.3

Figure 14-9: Discover Devices

18. Wait a minute, then refresh the browser. You will have a device entry.

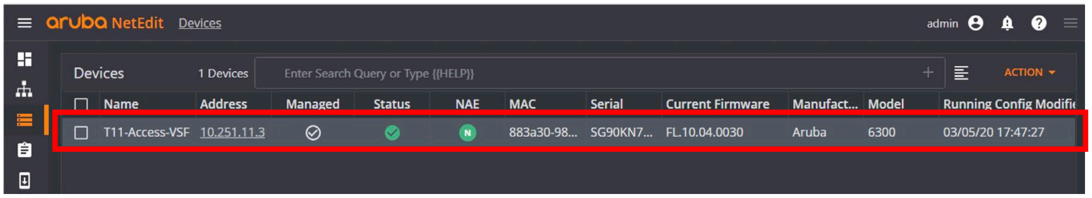

□	Name	Address	Managed	Status	NAE	MAC	Serial	Current Firmware	Manufact...	Model	Running Config Modifie
□	T11-Access-VSF	10.251.11.3	✓	✓	N	883a30-98...	SG90KN7...	FL.10.04.0030	Aruba	6300	03/05/20 17:47:27

Figure 14-10: NetEdit Devices 2

What device has NetEdit discovered?

____

What version is listed under Current Firmware column?

____

What model is it?

____

19. Click on the IP address of **Access-VSF**. That will take you to the Device Details page.

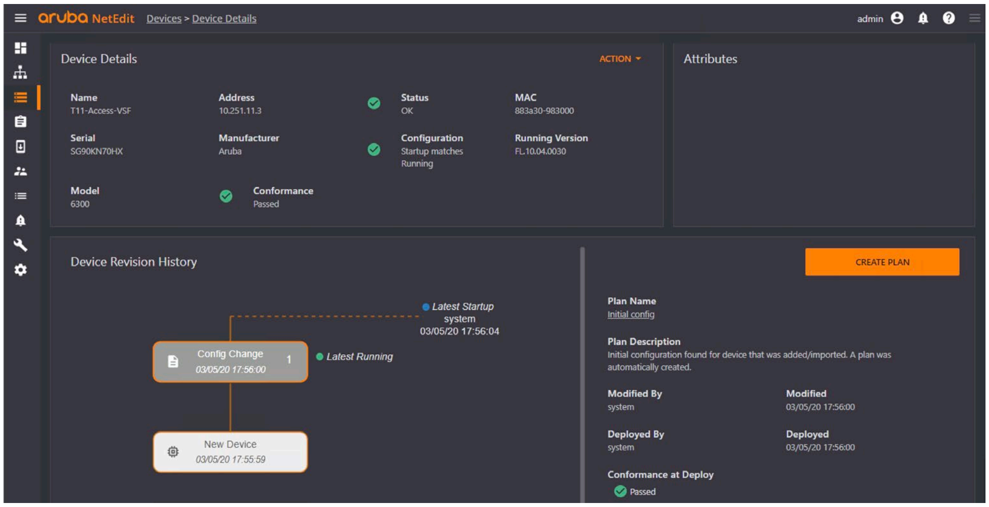
aruba NetEdit
Devices > Device Details
admin
ACTION
Attributes
CREATE PLAN
Config Change 03/05/20 17:56:00 1
New Device 03/05/20 17:55:59
Latest Startup system 03/05/20 17:56:04
Latest Running

Device Details

Name
T11-Access-VSF

Address
10.251.11.3

Status
OK

MAC
883a30-983000

Serial
SG90KN70HX

Manufacturer
Aruba

Configuration
Startup matches Running

Running Version
FL.10.04.0030

Model
6300

Conformance
Passed

Device Revision History

Plan Name
Initial config

Plan Description
Initial configuration found for device that was added/imported. A plan was automatically created.

Modified By
system

Modified
03/05/20 17:56:00

Deployed By
system

Deployed
03/05/20 17:56:00

Conformance at Deploy
✓ Passed

Figure 14-11: Device Details

NOTE: In addition to the regular details (Name, model, serial number, address and Code version) you can also see whether or not the Startup and Running configurations match, as well as their current Conformance and Status states.

20. Click the **Action** button, then select **Hardware Information** from the menu that appears. The Hardware Information dialog box will show up.
21. Expand the **Management Module** and **Power Supply** sections.

What Member of the VSF stack is considered the Management module?

____

How many power supplies are listed?

____

What physical device do they relate to?

____

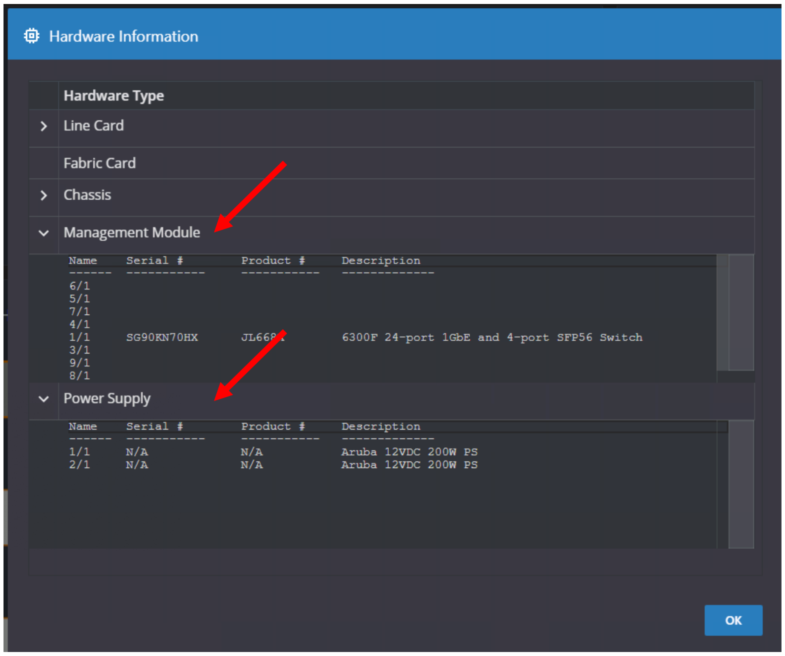

Management Module

```
Name    Serial #       Product #     Description
------  -----------    -----------   -------------
6/1
5/1
7/1
4/1
1/1     SG90KN70HX     JL668A        6300F 24-port 1GbE and 4-port SFP56 Switch
3/1
9/1
8/1
```

Power Supply

```
Name    Serial #       Product #     Description
------  -----------    -----------   -------------
1/1     N/A            N/A           Aruba 12VDC 200W PS
2/1     N/A            N/A           Aruba 12VDC 200W PS
```

Figure 14-12: Hardware Information

22. Click **OK** button.
23. Click the **Action** button, then select **View Firmware Information** from the menu that appears. The Firmware Information dialog box will show up.

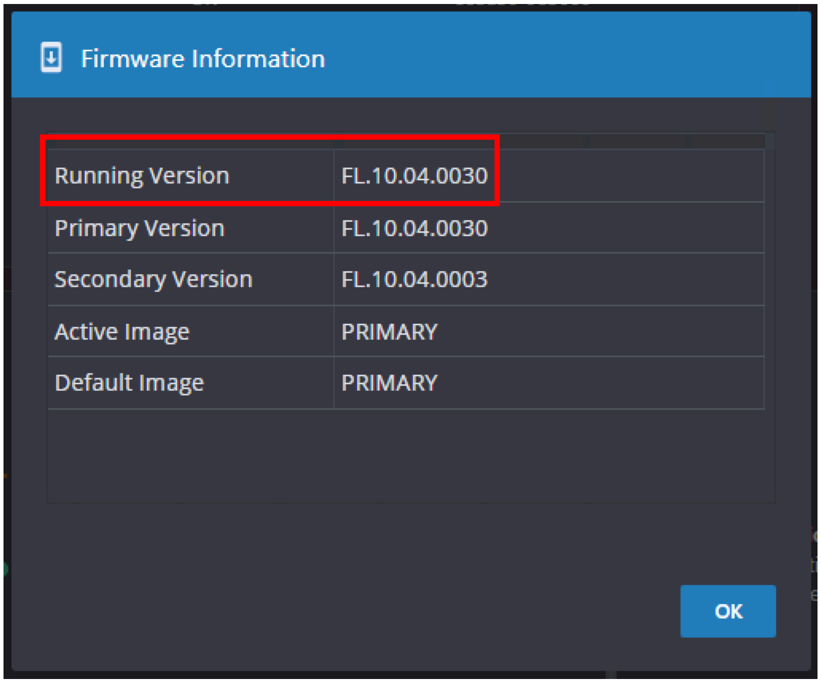

Firmware Information

Running Version	FL.10.04.0030
Primary Version	FL.10.04.0030
Secondary Version	FL.10.04.0003
Active Image	PRIMARY
Default Image	PRIMARY

Figure 14-13: Firmware Information

What version of code is running in the system?

___

What AOS-CX code versions are stored on Primary and Secondary partitions?

___

24. Click **OK** button.
25. Click the **Action** button, then select **View Running Config** from the menu that appears. The “Device Viewer RUNNING” section for Access-VSF shows up and will display the running configuration.

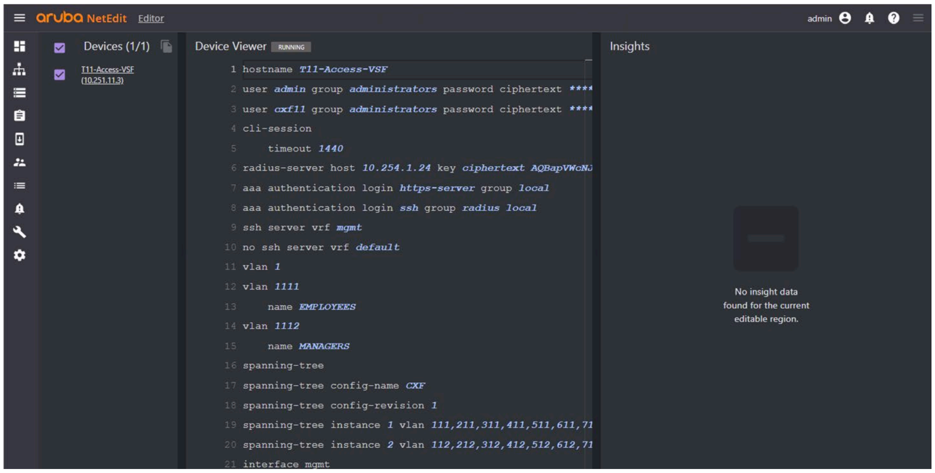

```
hostname T11-Access-VSF
user admin group administrators password ciphertext ****
user cxf11 group administrators password ciphertext ****
cli-session
    timeout 1440
radius-server host 10.254.1.24 key ciphertext AQBapVWcNJ
aaa authentication login https-server group local
aaa authentication login ssh group radius local
ssh server vrf mgmt
no ssh server vrf default
vlan 1
vlan 1111
    name EMPLOYEES
vlan 1112
    name MANAGERS
spanning-tree
spanning-tree config-name CXF
spanning-tree config-revision 1
spanning-tree instance 1 vlan 111,211,311,411,511,611,71
spanning-tree instance 2 vlan 112,212,312,412,512,612,71
interface mgmt
```

Figure 14-14: Device View

26. Click the **Overview** button (▦)in the navigation panel.

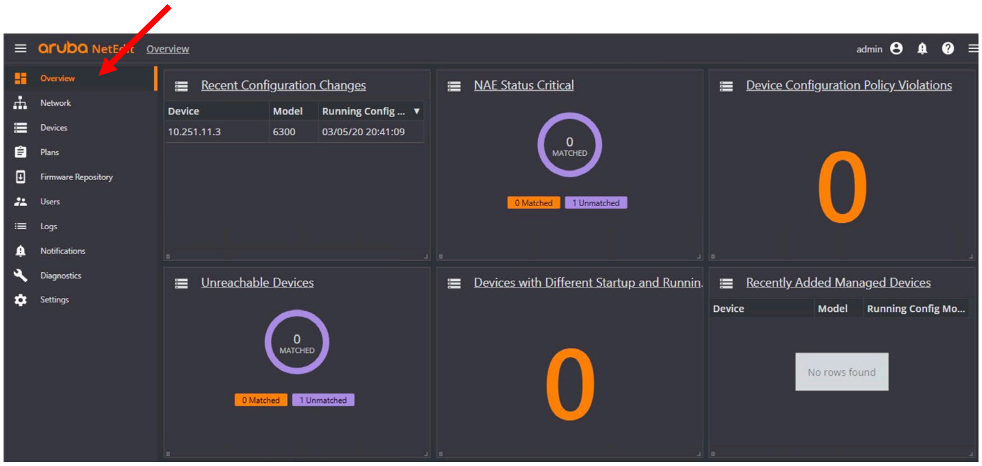

Device	Model	Running Config ... ▼
10.251.11.3	6300	03/05/20 20:41:09

Device	Model	Running Config Mo...

Figure 14-15: NetEdit Overview 2

Are there any unreachable devices?

____

Do any devices have different Startup and Running Configurations? If so, what are they?

____

Task 2: Deployment Plan

Objectives

The IT Staff at BigStartup seem quite impressed by NetEdit and its monitoring capabilities. However, they are wondering if configuration changes can be made from the tool. Next you will demonstrate NetEdit's script deployments capabilities as well as roll back options in case of configuration mistakes.

In this task you will run a deployment plan and commit it, so the configuration changes remain even if the devices reboots. Then you will inspect the NetEdit logs.

Steps

Access-VSF

1. Move to **Access-VSF**'s console.
2. Display the brief information of **port 2/1/4**.

```
T11-Access-VSF# show interface 2/1/4 brief
-------------------------------------------------------------------------------------
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
          VLAN                                                                 (Mb/s)
-------------------------------------------------------------------------------------
2/1/4     1111    access 1GbT           yes     up                             1000
T11-Access-VSF#
```

What VLAN is the port mapped to?

____

PC-1

3. Access **PC-1**.
4. Open a browser and type the NetEdit IP address in the URL field (**10.251.X.200**) then hit **[Enter]**.

5. Login with **admin** as username and **aruba123** as password.
6. Click on **Devices** in the navigation pane (☰).

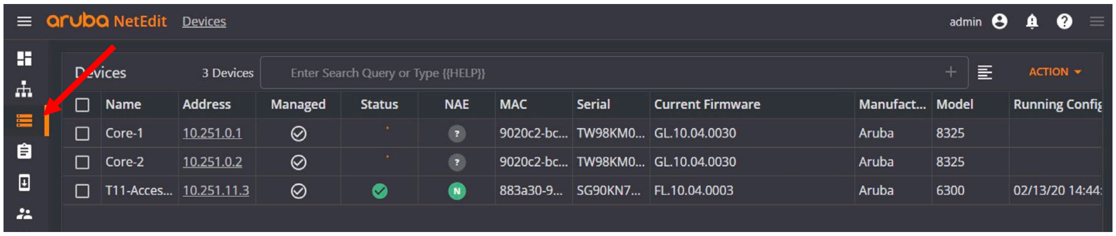

Devices 3 Devices Enter Search Query or Type {{HELP}}

□	Name	Address	Managed	Status	NAE	MAC	Serial	Current Firmware	Manufact...	Model	Running Config
□	Core-1	10.251.0.1	✓		?	9020c2-bc...	TW98KM0...	GL.10.04.0030	Aruba	8325	
□	Core-2	10.251.0.2	✓		?	9020c2-bc...	TW98KM0...	GL.10.04.0030	Aruba	8325	
□	T11-Acces...	10.251.11.3	✓	✓	N	883a30-9...	SG90KN7...	FL.10.04.0003	Aruba	6300	02/13/20 14:44:

Figure 14-16: NetEdit Devices 3

7. Check the checkbox next to **Access-VSF**’s IP address (**10.251.X.3**).
8. Click **ACTION** at far right, then select **Edit Config** from the menu that appears. This takes you to PLAN section and shows a Create Plan dialog box.

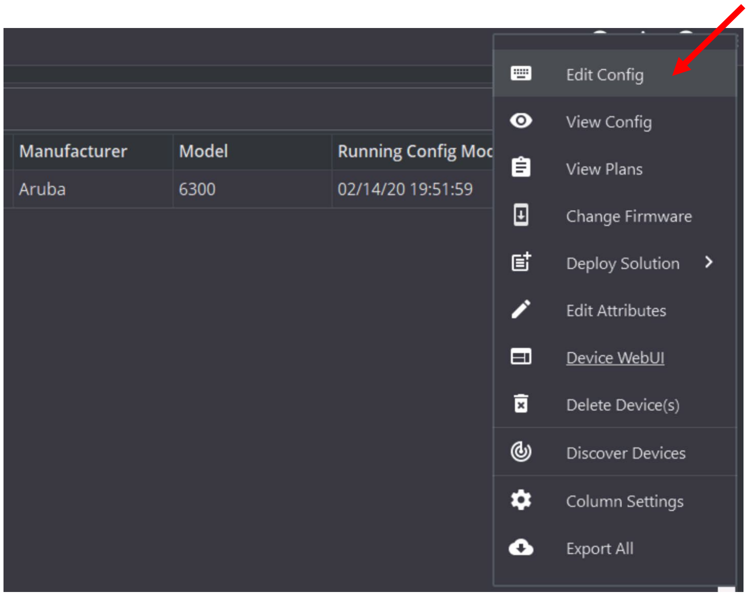

Manufacturer	Model	Running Config Moc
Aruba	6300	02/14/20 19:51:59

Figure 14-17: Actions Edit Config

9. Under name type **VLANX12**.
10. Under Description type **Assign VLANX12 to port 2/1/4**.

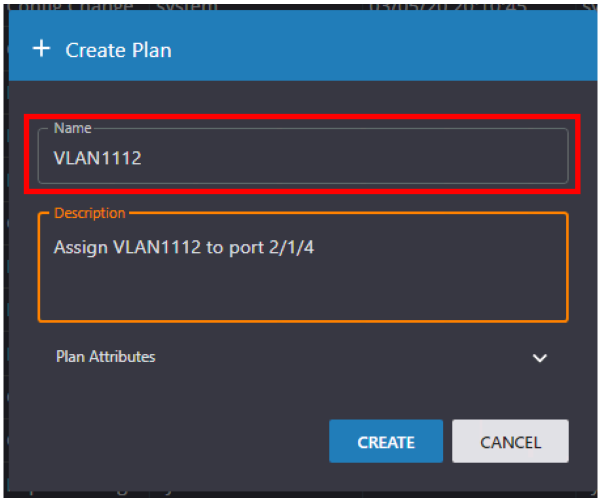

Name
VLAN1112

Description
Assign VLAN1112 to port 2/1/4

Figure 14-18: Create Plan

11. Click on the **CREATE** button.
12. In the configuration section, scroll down to **interface 2/1/4**.
13. Click on **vlan access X11** and change it to **vlan access X12**.

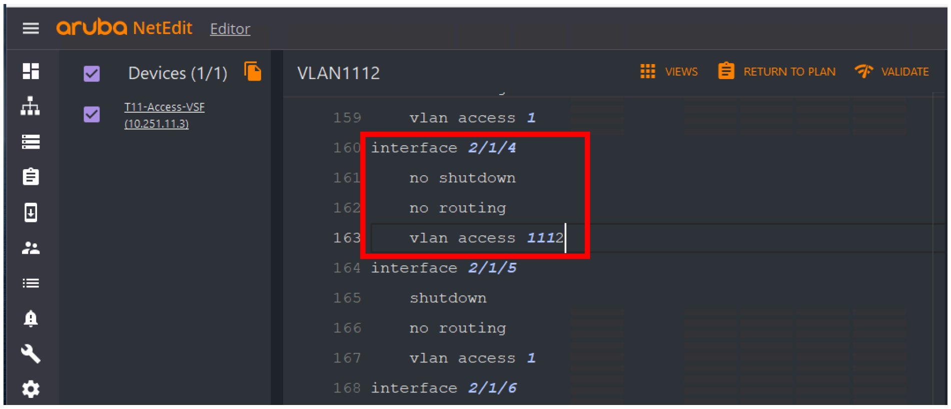

```
    vlan access 1
interface 2/1/4
    no shutdown
    no routing
    vlan access 1112
interface 2/1/5
    shutdown
    no routing
    vlan access 1
interface 2/1/6
```

Figure 14-19: VLAN X12 Configuration Plan

14. Click on **VALIDATE**. The plan will be validated and should be successful.

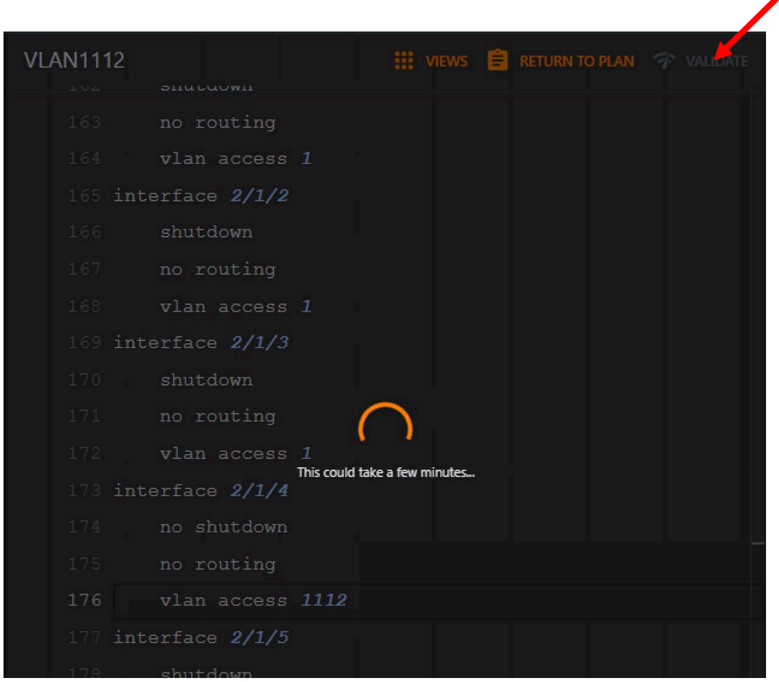

```
    shutdown
    no routing
    vlan access 1
interface 2/1/2
    shutdown
    no routing
    vlan access 1
interface 2/1/3
    shutdown
    no routing
    vlan access 1
interface 2/1/4
    no shutdown
    no routing
    vlan access 1112
interface 2/1/5
    shutdown
```

This could take a few minutes...

Figure 14-20: Plan Validation

15. Click on **RETURN TO PLAN**. This takes you to “*Plans > Plans Details*”.

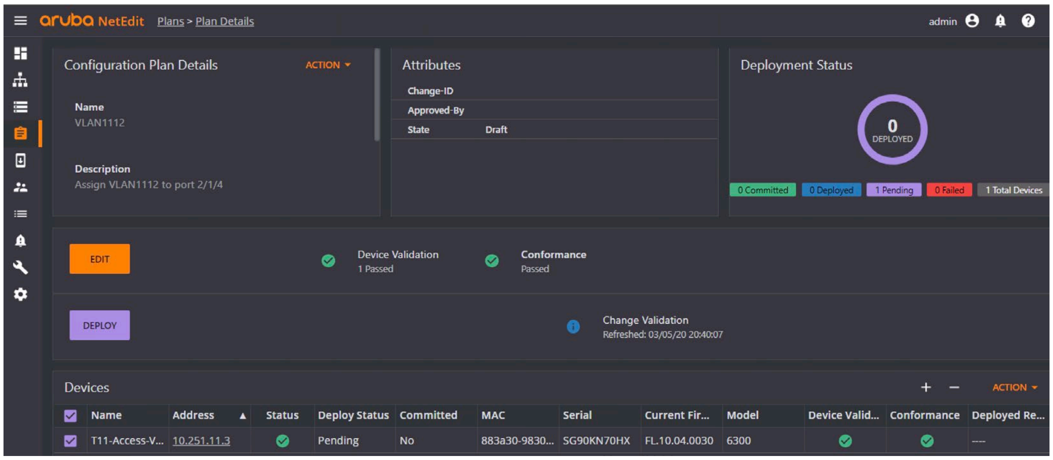

Configuration Plan Details

Name
VLAN1112

Description
Assign VLAN1112 to port 2/1/4

Attributes

Change-ID	
Approved-By	
State	Draft

Deployment Status

0 DEPLOYED

0 Committed | 0 Deployed | 1 Pending | 0 Failed | 1 Total Devices

✓ Device Validation
1 Passed

✓ Conformance
Passed

Change Validation
Refreshed: 03/05/20 20:40:07

Devices

☑	Name	Address ▲	Status	Deploy Status	Committed	MAC	Serial	Current Fir...	Model	Device Valid...	Conformance	Deployed Re...
☑	T11-Access-V...	10.251.11.3	✓	Pending	No	883a30-9830...	SG90KN70HX	FL.10.04.0030	6300	✓	✓	----

Figure 14-21: Plan Details

16. Confirm that your newly created plan is listed. Then click the purple **DEPLOY** button. A dialog box appears.

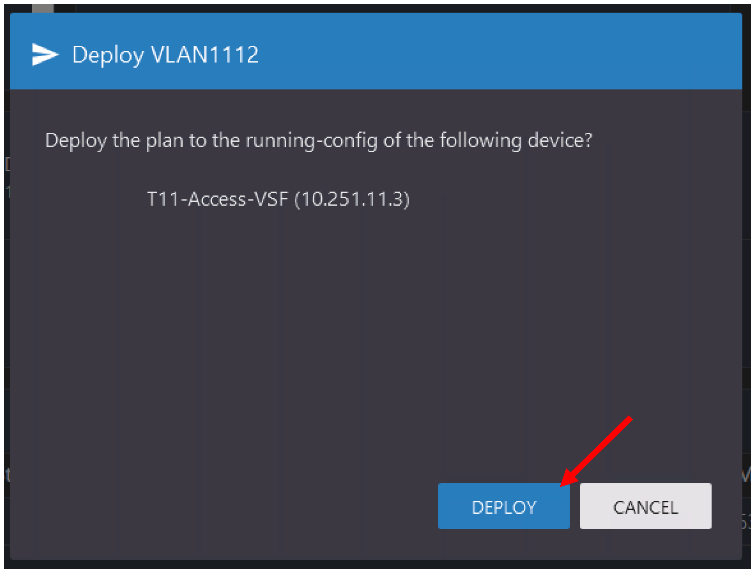

Deploy the plan to the running-config of the following device?

T11-Access-VSF (10.251.11.3)

Figure 14-22: Deploy VLAN X12

17. Click the blue **DEPLOY** button, you should receive a "*Deployment is in progress…"* message at the bottom right.

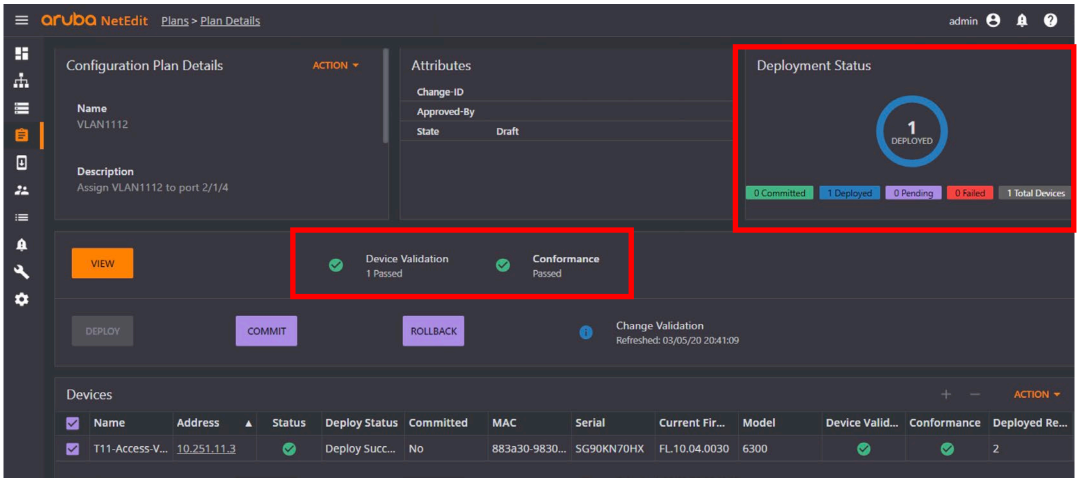

Configuration Plan Details

Name
VLAN1112

Description
Assign VLAN1112 to port 2/1/4

Attributes

Change-ID	
Approved-By	
State	Draft

Deployment Status

1 DEPLOYED

0 Committed | 1 Deployed | 0 Pending | 0 Failed | 1 Total Devices

Device Validation
1 Passed

Conformance
Passed

Change Validation
Refreshed: 03/05/20 20:41:09

Devices

☑	Name	Address ▲	Status	Deploy Status	Committed	MAC	Serial	Current Fir...	Model	Device Valid...	Conformance	Deployed Re...
☑	T11-Access-V...	10.251.11.3	✓	Deploy Succ...	No	883a30-9830...	SG90KN70HX	FL.10.04.0030	6300	✓	✓	2

Figure 14-23: Deployment is in Progress

What was the Device Validation Result?

____

What was the Conformance Result?

What is the Deployment Status?

18. Click on **COMMIT**. A dialog box appears.

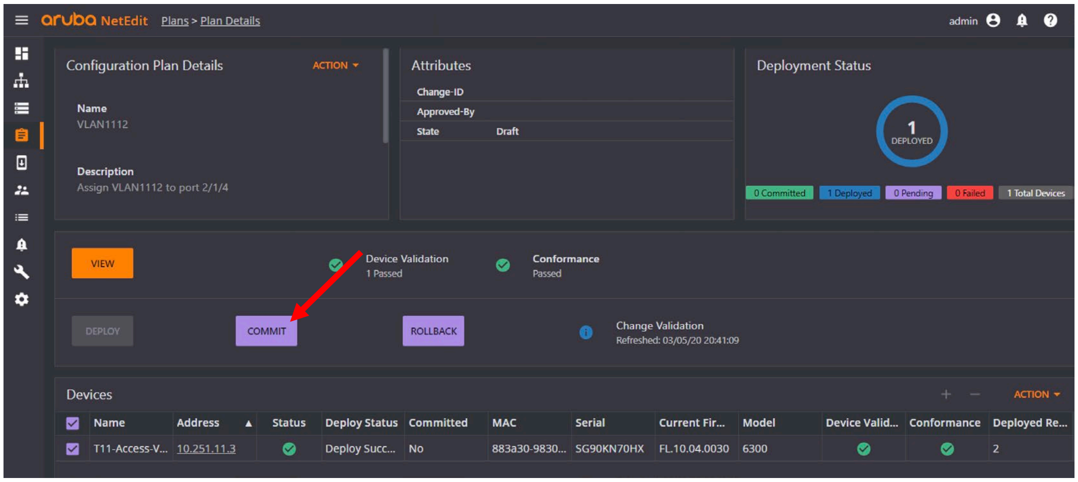

Configuration Plan Details

Name
VLAN1112

Description
Assign VLAN1112 to port 2/1/4

Attributes

Change-ID	
Approved-By	
State	Draft

Deployment Status

1 DEPLOYED

0 Committed | 1 Deployed | 0 Pending | 0 Failed | 1 Total Devices

Device Validation
1 Passed

Conformance
Passed

Change Validation
Refreshed: 03/05/20 20:41:09

Devices

☑	Name	Address ▲	Status	Deploy Status	Committed	MAC	Serial	Current Fir...	Model	Device Valid...	Conformance	Deployed Re...
☑	T11-Access-V...	10.251.11.3	✓	Deploy Succ...	No	883a30-9830...	SG90KN70HX	FL.10.04.0030	6300	✓	✓	2

Figure 14-24: Deployed plan

19. Click on **COMMIT** again. This will save the current configuration in the "startup-config" checkpoint.

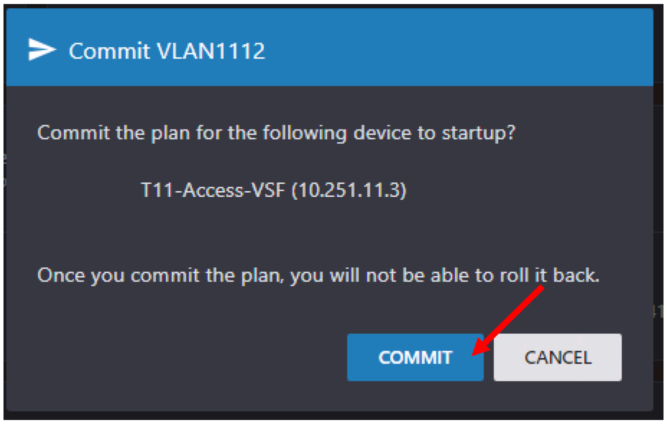

Commit VLAN1112

Commit the plan for the following device to startup?

T11-Access-VSF (10.251.11.3)

Once you commit the plan, you will not be able to roll it back.

Figure 14-25: Committing a plan

What is the Deployment Status now?

____

Access-VSF

20. Move back to **Access-VSF**'s console.
21. Display the brief information of **port 2/1/4**.

```
T11-Access-VSF# show interface 2/1/4 brief
-------------------------------------------------------------------------------------
Port       Native  Mode    Type           Enabled Status  Reason                Speed
           VLAN                                                                 (Mb/s)
-------------------------------------------------------------------------------------
2/1/4      1112    access 1GbT            yes     up                            1000
T11-Access-VSF#
```

What VLAN is the port mapped to now?

____

PC-1

22. Move back to **PC-1**.
23. Click on **Logs** in the left navigation pane (☰). You will see evidence of the previous deployment.

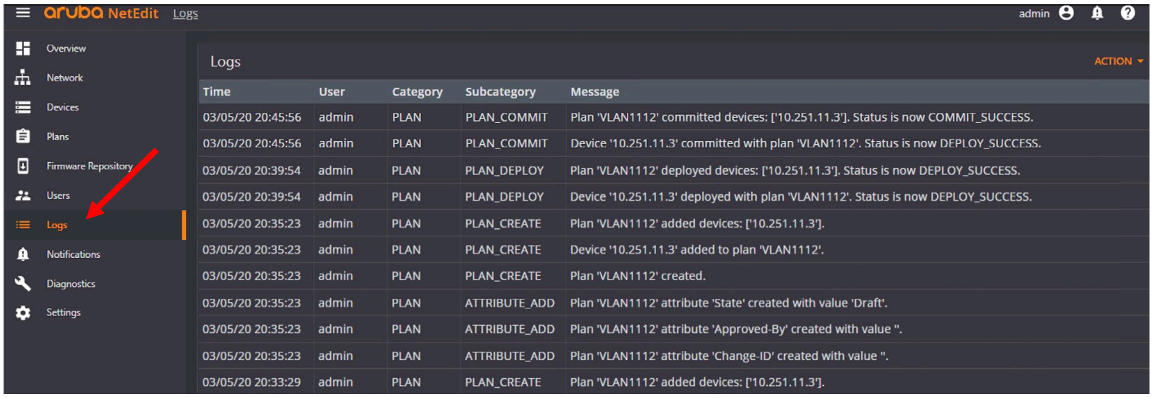

Logs

Time	User	Category	Subcategory	Message
03/05/20 20:45:56	admin	PLAN	PLAN_COMMIT	Plan 'VLAN1112' committed devices: ['10.251.11.3']. Status is now COMMIT_SUCCESS.
03/05/20 20:45:56	admin	PLAN	PLAN_COMMIT	Device '10.251.11.3' committed with plan 'VLAN1112'. Status is now DEPLOY_SUCCESS.
03/05/20 20:39:54	admin	PLAN	PLAN_DEPLOY	Plan 'VLAN1112' deployed devices: ['10.251.11.3']. Status is now DEPLOY_SUCCESS.
03/05/20 20:39:54	admin	PLAN	PLAN_DEPLOY	Device '10.251.11.3' deployed with plan 'VLAN1112'. Status is now DEPLOY_SUCCESS.
03/05/20 20:35:23	admin	PLAN	PLAN_CREATE	Plan 'VLAN1112' added devices: ['10.251.11.3'].
03/05/20 20:35:23	admin	PLAN	PLAN_CREATE	Device '10.251.11.3' added to plan 'VLAN1112'.
03/05/20 20:35:23	admin	PLAN	PLAN_CREATE	Plan 'VLAN1112' created.
03/05/20 20:35:23	admin	PLAN	ATTRIBUTE_ADD	Plan 'VLAN1112' attribute 'State' created with value 'Draft'.
03/05/20 20:35:23	admin	PLAN	ATTRIBUTE_ADD	Plan 'VLAN1112' attribute 'Approved-By' created with value ''.
03/05/20 20:35:23	admin	PLAN	ATTRIBUTE_ADD	Plan 'VLAN1112' attribute 'Change-ID' created with value ''.
03/05/20 20:33:29	admin	PLAN	PLAN_CREATE	Plan 'VLAN1112' added devices: ['10.251.11.3'].

Figure 14-26: NetEdit Logs

You have completed Lab 14!

AOS-CX Switching Fundamentals

Appendix 1: Numerical conversion

Overview

This appendix contains the results of Lab 1 exercises.

Task 1: Binary to Decimal conversion

Objectives

Convert the following binary into decimal values.

a) 10101010
b) 11100011
c) 01110000 (optional)
d) 10000001 (optional)
e) 00011100 (optional)

Results

Table 1-1: Power of 2: Binary to decimal

Powers of 2	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
Decimal	**128**	**64**	**32**	**16**	**8**	**4**	**2**	**1**

Binary a)	**1**	**0**	**1**	**0**	**1**	**0**	**1**	**0**
Decimal a)	**128**	**0**	**32**	**0**	**8**	**0**	**2**	**0**

128+32 +8+2=**170**

Binary b)	**1**	**1**	**1**	**0**	**0**	**0**	**1**	**1**
Decimal b)	**128**	**64**	**32**	**0**	**0**	**0**	**2**	**1**

128+64+32+2+1= **227**

Binary c)	**0**	**1**	**1**	**1**	**0**	**0**	**0**	**0**
Decimal c)	**0**	**64**	**32**	**16**	**0**	**0**	**0**	**0**

64+32+16 = **112**

Binary d)	1	0	0	0	0	0	0	1
Decimal d)	128	0	0	0	0	0	0	1

128+1 = **129**

Binary e)	0	0	0	1	1	1	0	0
Decimal e)	0	0	0	16	8	4	0	0

16+8+4 = **28**

Task 2: Decimal to Binary conversion method 1

Objectives

Convert the following decimal values into binary using the division method:

a) 315
b) 116
c) 39 (optional)
d) 240 (optional)

Results

a) Convert 315

Rem. 9	Rem. 8	Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1
1	0	0	1	1	1	0	1	1

b) Convert 116

Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1
1	1	1	0	1	0	0

c) Convert 39

Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1
1	0	0	1	1	1

d) Convert 240

Rem. 8	Rem. 7	Rem. 6	Rem. 5	Rem. 4	Rem. 3	Rem. 2	Rem. 1
1	1	1	1	0	0	0	0

Task 3: Decimal to Binary conversion method 2

Objectives

Convert the following decimal values into binary using the division method.

a) 224
b) 17
c) 199 (optional)
d) 46 (optional)

Results

Table 1-2: Power of 2: Decimal to binary

Powers of 2	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
Decimal	128	64	32	16	8	4	2	1
Binary a)	1	1	1	0	0	0	0	0
Binary b)	0	0	0	1	0	0	0	1
Binary c)	1	1	0	0	0	1	1	1
Binary d)	0	0	1	0	1	1	1	0

Task 4: Binary to Hexadecimal conversion

Objectives

Convert the following binary values into hexadecimal.

a) 01100110
b) 10100101
c) 00010010 (optional)
d) 01011010 (optional)

Results

Table 1-3: Binary to Hexadecimal

Binary	Hexadecimal
0000	0x0
0001	0x1
0010	0x2
0011	0x3
0100	0x4
0101	0x5
0110	0x6
0111	0x7
1000	0x8
1001	0x9
1010	0xA
1011	0xB
1100	0xC
1101	0xD
1110	0xE
1111	0xF

a) Convert 01100110 = **66**

b) Convert 10100101 = **A5**

c) Convert 00010010 = **12**

d) Convert 01011010 = **5A**

Task 5: Hexadecimal to Binary conversion

Objectives

Convert the following decimal values into binary using the division method.

a) AB
b) AB3
c) 3F4 (optional)
d) 0C (optional)

Results

a) Convert AB = **10101011**

b) Convert AB3 = **101010110011**

c) Convert 3F4 = **001111110100**

d) Convert 0C = **00001100**

Task 6: Decimal to Hexadecimal conversion (optional)

Objectives

Convert the following decimal values into binary using the division method.

a) 898
b) 2033
c) 1572
d) 78

Results

Table 1-4: Decimal to Hexadecimal

Decimal	Hexadecimal
0	0x0
1	0x1
2	0x2
3	0x3
4	0x4
5	0x5
6	0x6
7	0x7
8	0x8
9	0x9
10	0xA
11	0xB
12	0xC
13	0xD
14	0xE
15	0xF

a) Convert 898

Rem. 3	Rem. 2	Rem. 1
3	8	2

b) Convert 2033

Rem. 3	Rem. 2	Rem. 1
7	F	1

c) Convert 1572

Rem. 3	Rem. 2	Rem. 1
6	2	4

d) Convert 78

Rem. 2	Rem. 1
4	E

Task 7: Hexadecimal to Decimal conversion (optional)

Objectives

Convert the following decimal values into binary using the division method.

a) F3A
b) 15B
c) 111
d) 7C

Results

Table 1-5: Decimal to Hexadecimal

Power of 16	16^3	16^2	16^2	16^0
Decimal	**4096**	**256**	**16**	**1**

Hexadecimal a)		F	3	A
Decimal		15	3	10
Multiplication		3840	48	10
Decimal a)	3840+48+10=**3898**			

Hexadecimal b)		1	5	B
Decimal		1	5	11
Multiplication		256	80	11
Decimal b)	256 + 80 + 11 = **347**			

Hexadecimal c)		1	1	1

Decimal		1	1	1
Multiplication		256	16	1
Decimal c)	256+16+1 = **273**			

Hexadecimal d)			7	C
Decimal			7	12
Multiplication			112	12
Decimal d)	112+12 = **124**			

AOS-CX Switching Fundamentals

Appendix 2: Subnetting and VLSM

Overview

This appendix contains the results of Lab 9 exercises.

Task 1: Class A Subnetting

Objectives

Subnet the prefix using the information below:

Network Address: **43.0.0.0**
Number of needed Subnets: **9**

Steps

1. List all subnets in table 9-1 down below.

   What is the address class?

   Class A

   What is the default subnet mask?

   255.0.0.0

   What is the required subnet mask?

   255.240.0.0

   How many subnets will be generated with equal length subnet mask?

   2^S = 2^4 = **16 subnets**

   What is the total number of assignable addresses per subnet?

   2^H-2 = 2^{20}-2 = 1,048,576-2 = **1,048,574 addresses**

   How many bits were borrowed from the host portion in the default mask for creating subnets?

   4 bits

Table 9-1: Subnetting Task 1

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	43.0.0.0	43.0.0.1	43.15.255.254	43.15.255.255
2	43.16.0.0	43.16.0.1	43.31.255.254	43.31.255.255
3	43.32.0.0	43.32.0.1	43.47.255.254	43.47.255.255
4	43.48.0.0	43.48.0.1	43.63.255.254	43.63.255.255
5	43.64.0.0	43.64.0.1	43.79.255.254	43.79.255.255
6	43.80.0.0	43.80.0.1	43.95.255.254	43.95.255.255
7	43.96.0.0	43.96.0.1	43.111.255.254	43.111.255.255
8	43.112.0.0	43.112.0.1	43.127.255.254	43.127.255.255
9	43.128.0.0	43.128.0.1	43.143.255.254	43.143.255.255
10	43.144.0.0	43.144.0.1	43.159.255.254	43.159.255.255
11	43.160.0.0	43.160.0.1	43.175.255.254	43.175.255.255
12	43.176.0.0	43.176.0.1	43.191.255.254	43.191.255.255
13	43.192.0.0	43.192.0.1	43.207.255.254	43.207.255.255
14	43.208.0.0	43.208.0.1	43.223.255.254	43.223.255.255
15	43.224.0.0	43.224.0.1	43.239.255.254	43.239.255.255
16	43.240.0.0	43.240.0.1	43.255.255.254	43.255.255.255

Task 2: Class B Subnetting

Objectives

Subnet the prefix using the information below:

IP Address: **132.89.5.10**

Number of needed Subnets: **20**

Steps

1. List all subnets in table 9-2 down below

What network does the address belong to?

132.89.0.0/16

What is the address class?

Class B

What is the default subnet mask?

255.255.0.0

What is the required subnet mask?

255.255.248.0

How many subnets will be generated with equal length subnet mask?

2^S = 2^5 = **32 subnets**

How many bits were borrowed from the host portion in the default mask for creating subnets?

5 bits

Table 9-2: Subnetting Task 2

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	132.89.0.0	132.89.0.1	132.89.7.254	132.89.7.255
2	132.89.8.0	132.89.8.1	132.89.15.254	132.89.15.255
3	132.89.16.0	132.89.16.1	132.89.23.254	132.89.23.255
4	132.89.24.0	132.89.24.1	132.89.31.254	132.89.31.255
5	132.89.32.0	132.89.32.1	132.89.39.254	132.89.39.255
6	132.89.40.0	132.89.40.1	132.89.47.254	132.89.47.255
7	132.89.48.0	132.89.48.1	132.89.55.254	132.89.55.255
8	132.89.56.0	132.89.56.1	132.89.63.254	132.89.63.255
9	132.89.64.0	132.89.64.1	132.89.71.254	132.89.71.255
10	132.89.72.0	132.89.72.1	132.89.79.254	132.89.79.255
11	132.89.80.0	132.89.80.1	132.89.87.254	132.89.87.255
12	132.89.88.0	132.89.88.1	132.89.95.254	132.89.95.255
13	132.89.96.0	132.89.96.1	132.89.103.254	132.89.103.255
14	132.89.104.0	132.89.104.1	132.89.111.254	132.89.111.255
15	132.89.112.0	132.89.112.1	132.89.119.254	132.89.119.255
16	132.89.120.0	132.89.120.1	132.89.127.254	132.89.127.255
17	132.89.128.0	132.89.128.1	132.89.135.254	132.89.135.255
18	132.89.136.0	132.89.136.1	132.89.143.254	132.89.143.255
19	132.89.144.0	132.89.144.1	132.89.151.254	132.89.151.255
20	132.89.152.0	132.89.152.1	132.89.159.254	132.89.159.255
21	132.89.160.0	132.89.160.1	132.89.167.254	132.89.167.255
22	132.89.168.0	132.89.168.1	132.89.175.254	132.89.175.255
23	132.89.176.0	132.89.176.1	132.89.183.254	132.89.183.255
24	132.89.184.0	132.89.184.1	132.89.191.254	132.89.191.255
25	132.89.192.0	132.89.192.1	132.89.199.254	132.89.199.255
26	132.89.200.0	132.89.200.1	132.89.207.254	132.89.207.255
27	132.89.208.0	132.89.208.1	132.89.215.254	132.89.215.255
28	132.89.216.0	132.89.216.1	132.89.223.254	132.89.223.255
29	132.89.224.0	132.89.224.1	132.89.231.254	132.89.231.255

30	**132.89.232.0**	**132.89.232.1**	**132.89.239.254**	**132.89.239.255**
31	**132.89.240.0**	**132.89.240.1**	**132.89.247.254**	**132.89.247.255**
32	**132.89.248.0**	**132.89.248.1**	**132.89.255.254**	**132.89.255.255**

Task 3a: Class C Subnetting Part 1

Objectives

Subnet the prefix using the information below:

Network Address: **192.168.1.0**
Number of needed assignable host addresses: **2**

Steps

1. List the first 4 subnets and the last one in table 9-3

What is the address class?

Class C

What is the default subnet mask?

255.255.255.0

What is the required subnet mask?

255.255.255.252

How many subnets will be generated with equal length subnet mask?

$2^S = 2^6 =$ **64 subnets**

What is the total number of assignable addresses per subnet?

$2^H-2 = 2^2-2 = 4-2 =$ **2 addresses**

How many bits were borrowed from the host portion in the default mask for creating subnets?

6 bits

Table 9-3: Subnetting Task 3a

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	192.168.1.0	192.168.1.1	192.168.1.3	192.168.1.3
2	192.168.1.4	192.168.1.5	192.168.1.6	192.168.1.7
3	192.168.1.8	192.168.1.9	192.168.1.10	192.168.1.11
4	192.168.1.12	192.168.1.13	192.168.1.14	192.168.1.15
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
64	192.168.1.252	192.168.1.253	192.168.1.254	192.168.1.255

Task 3b: Class C Subnetting Part 2 (optional)

Objectives

Subnet the prefix using the information below:

Network Address: **199.209.0.0**
Number of needed Subnets: **7**

Steps

1. List all subnets in table 9-4 down below.

What is the address class?

Class C

What is the default subnet mask?

255.255.255.0

What is the required subnet mask?

255.255.255.224

How many subnets will be generated with equal length subnet mask?

$2^S = 2^3 =$ **8 subnets**

What is the total number of assignable addresses per subnet?

$2^H-2 = 2^5-2 = 32-2 =$ **30 addresses**

How many bits were borrowed from the host portion in the default mask for creating subnets?

3 bits

Table 9-4: Subnetting Task 3b

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	199.209.0.0	199.209.0.1	199.209.0.30	199.209.0.31
2	199.209.0.32	199.209.0.33	199.209.0.62	199.209.0.63
3	199.209.0.64	199.209.0.65	199.209.0.94	199.209.0.95
4	199.209.0.96	199.209.0.97	199.209.0.126	199.209.0.127
5	199.209.0.128	199.209.0.129	199.209.0.158	199.209.0.159
6	199.209.0.160	199.209.0.161	199.209.0.190	199.209.0.191
7	199.209.0.192	199.209.0.193	199.209.0.222	199.209.0.223
8	199.209.0.224	199.209.0.225	199.209.0.254	199.209.0.255

Task 4a: VLSM Prefixes

Objectives

Subnet the prefix using the information below:

Network Address: **10.0.0.0**
Number of needed assignable host addresses: **254**

Steps

1. List 1st, 2nd, 3rd,, 21th, 22th and the 101th subnets in table 9-5 down below

What is the address class?

Class A

What is the default subnet mask?

255.0.0.0

What is the required subnet mask?

255.255.255.0

How many subnets will be generated with equal length subnet mask?

2^S = 2^{16} = **65536 subnets**

What is the total number of assignable addresses per subnet?

2^H-2 = 2^8-2 = 256-2 = **254 addresses**

How many bits were borrowed from the host portion in the default mask for creating subnets?

16 bits

Table 9-5: Subnetting exercise 4a

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	10.0.0.0	10.0.0.1	10.0.0.254	10.0.0.255
2	10.0.1.0	10.0.1.1	10.0.1.254	10.0.1.255
3	10.0.2.0	10.0.2.1	10.0.2.254	10.0.2.255
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
21	10.0.20.0	10.0.20.1	10.0.20.254	10.0.20.255
22	10.0.21.0	10.0.21.1	10.0.21.254	10.0.21.255
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
101	10.0.100.0	10.0.100.1	10.0.100.254	10.0.100.255

Task 4b: VLSM – Point to Point Segments

Objectives

Subnet the prefix using the information below:

Take the first /24 subnet of exercise 4a and subnet it again with segments that support up to 2 assignable addresses.

Steps

1. List the first 5 subnets in table 9-6.

   What is the required subnet mask?

   255.255.255.252

   How many bits were borrowed from the host portion in the default mask for creating subnets?

   6 bits

Table 9-6: Subnetting exercise 4b

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	**10.0.0.0**	10.0.0.1	10.0.0.2	10.0.0.3
2	**10.0.0.4**	10.0.0.5	10.0.0.6	10.0.0.7
3	**10.0.0.8**	10.0.0.9	10.0.0.10	10.0.0.11
4	**10.0.0.12**	10.0.0.13	10.0.0.14	10.0.0.15
5	**10.0.0.16**	10.0.0.17	10.0.0.18	10.0.0.19

IMPORTANT: It is always a best practice to deploy a /30 prefix when the segment will be used on a link (physical or virtual) that only interconnects two Layer 3 devices e.g. Ethernet links between two routers or multilayer switches, GRE tunnels, serial links, etc.

Task 4c: VLSM – Grouping Two Subnets (optional)

Objectives

Subnet the prefix using the information below:

Combine subnets 21 and 22 of exercise 4a into a single one that supports 500 hosts.

TIP: Since 10.0.20.0/24 and 10.0.21.0/24 are contiguous networks you can combine them by using the Network ID of the first subnet and reduce its prefix length by 1.

Steps

1. List the resulting subnet in table 9-6.

What is the total number of assignable addresses?

$2^{H}-2 = 2^{9}-2 = 512-2 =$ **510 addresses**

What is the required subnet mask?

255.255.254.0

How many bits were borrowed from the host portion in the default mask for creating subnets?

15 bits were borrowed from the classful network subnet mask (255.0.0.0).

Table 9-6: Subnetting Task 4c

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	10.0.20.0	10.0.20.1	10.0.21.254	10.0.21.255

IMPORTANT: Multiple contiguous subnets can be combined together into a larger one in order to provide more assignable addresses within the same segment, or for summarization purposes when using static routes or dynamic routing protocols.

Task 4d: VLSM – Loopback Segments (optional)

Objectives

Subnet the prefix using the information below:

Use the 101th subnet of exercise 4a and subnet it again with segments that support up to 1 host address.

Steps

1. List the first 5 subnets in table 9-7.

What is the required subnet mask?

255.255.255.255

How many subnets will be generated with this new subnet mask out of the 10.0.100.0/24 one?

$2^S = 2^8 =$ **256 subnets**

How many bits were borrowed from the host portion in the custom 255.255.255.0 mask for creating these subnets?

8 bits

Table 9-7: Subnetting exercise 5d

Subnet #	Network Identifier	1st assignable address	Last assignable address	Broadcast Address
1	10.0.100.0	Same as Network Identifier	Same as Network Identifier	Same as Network Identifier
2	10.0.100.1			
3	10.0.100.2			
4	10.0.100.3			
5	10.0.100.4			

IMPORTANT: /32 prefixes with a single host address can be used on loopback interfaces for connectivity tests, management and routing protocols for fine tuning (OSPF Router ID reachability, iBGP and eBGP multihop peering, etc). The device that owns the address is the only one with direct access to it unless you tell other devices how to reach it with static or dynamic routing since the address will be its segment itself!

It is always a good practice to reserve a range of addresses of your IP address scheme for this purpose and allocate one of them to each Layer 3 device in the network.

AOS-CX Switching Fundamentals

Appendix 3: Lab Diagrams

Overview

This appendix contains the Lab diagrams. Feel free to print them out and write on them when needed.

AOS-CX Switching Fundamentals Lab Topology

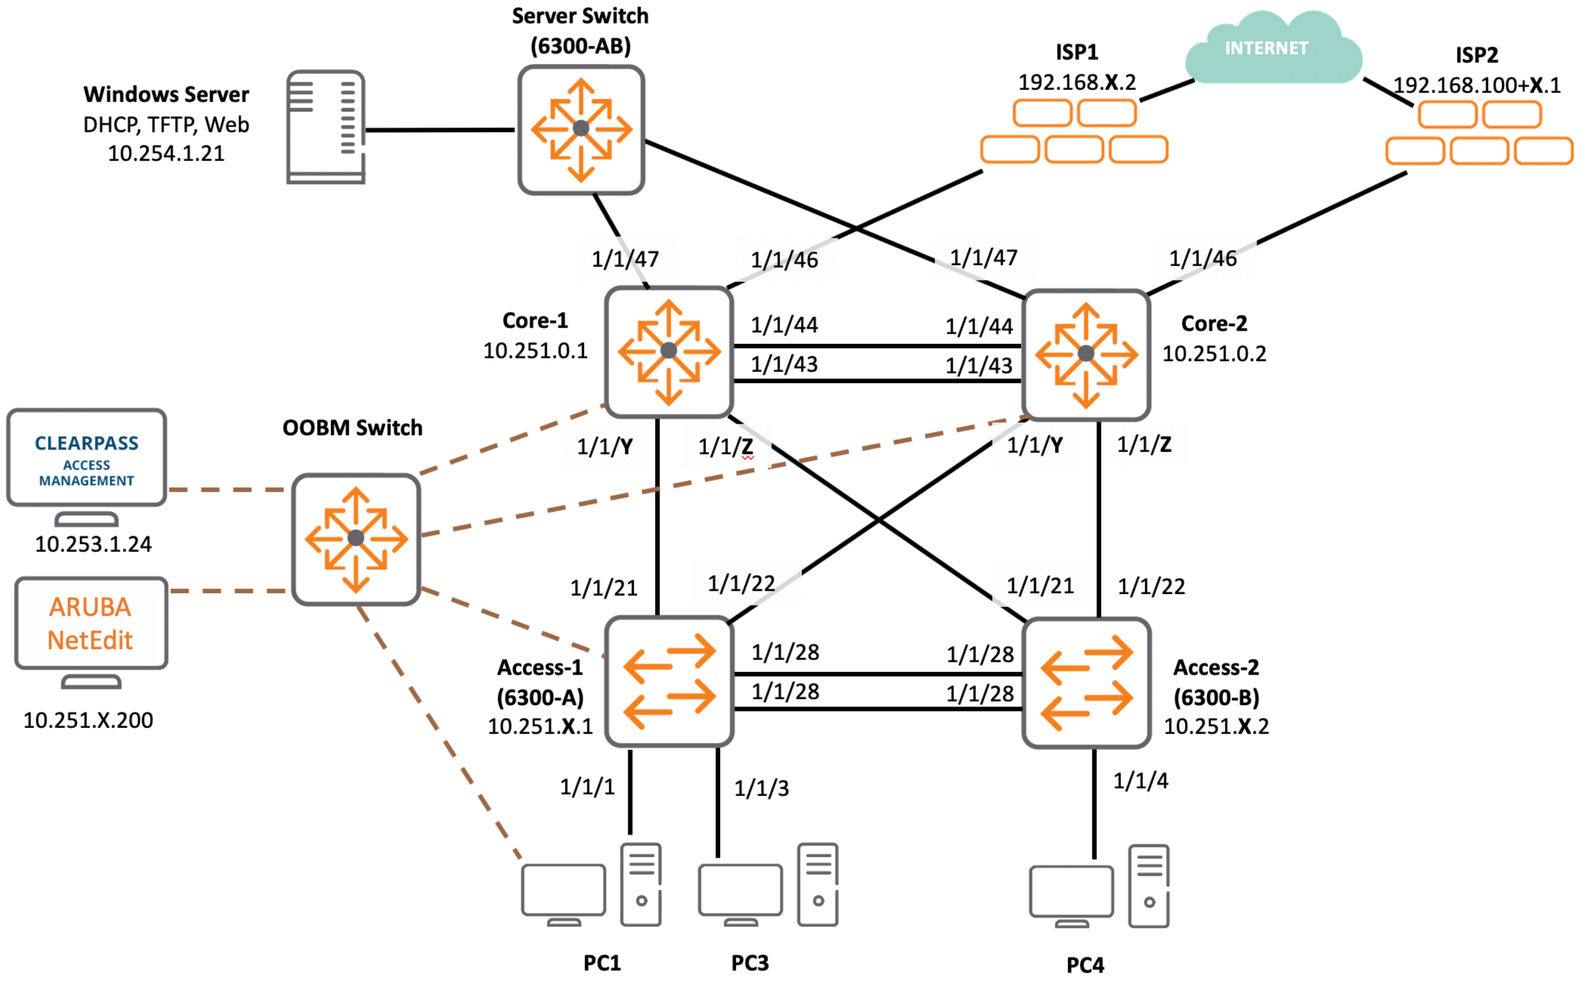

Mac Address	Port	VLAN
PC-1's MAC		
PC-3's MAC		

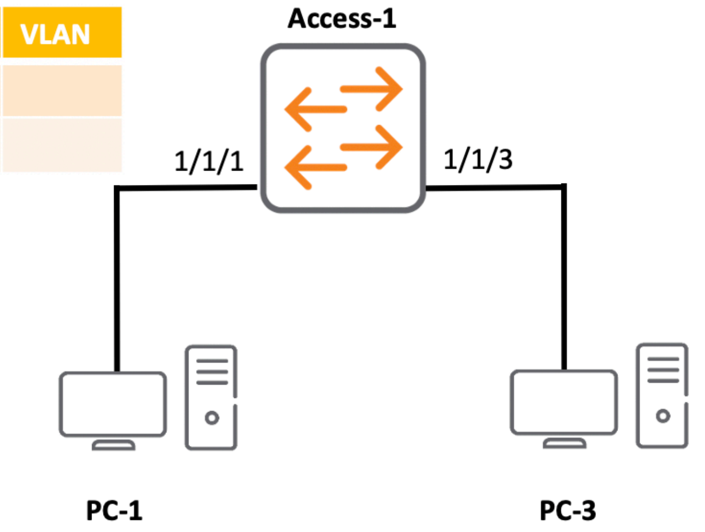

MAC:____________________ MAC:____________________

AOS-CX Switching Fundamentals

Figure 4.1-12: Access-1's MAC Address Table

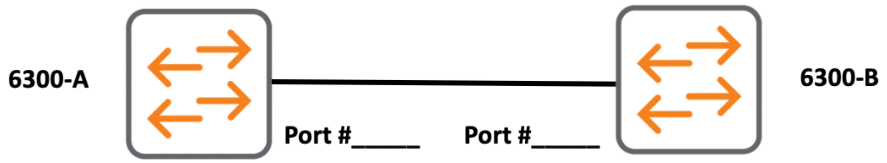

System Name:____________________

Chassis ID: ______________________

System Name:____________________

Chassis ID: ______________________

AOS-CX Switching Fundamentals

Figure 4.2-9: LLDP Discovery

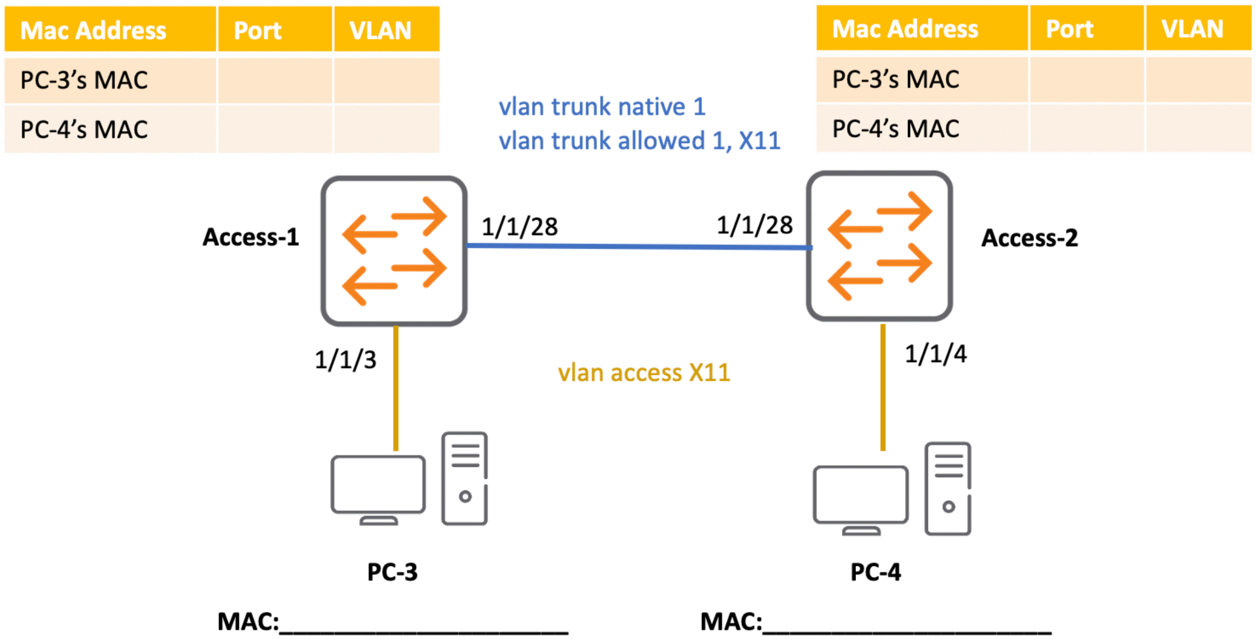

Mac Address	Port	VLAN
PC-3's MAC		
PC-4's MAC		

Mac Address	Port	VLAN
PC-3's MAC		
PC-4's MAC		

MAC:______________________

MAC:______________________

AOS-CX Switching Fundamentals
Figure 4.2-13: MAC Address Tables

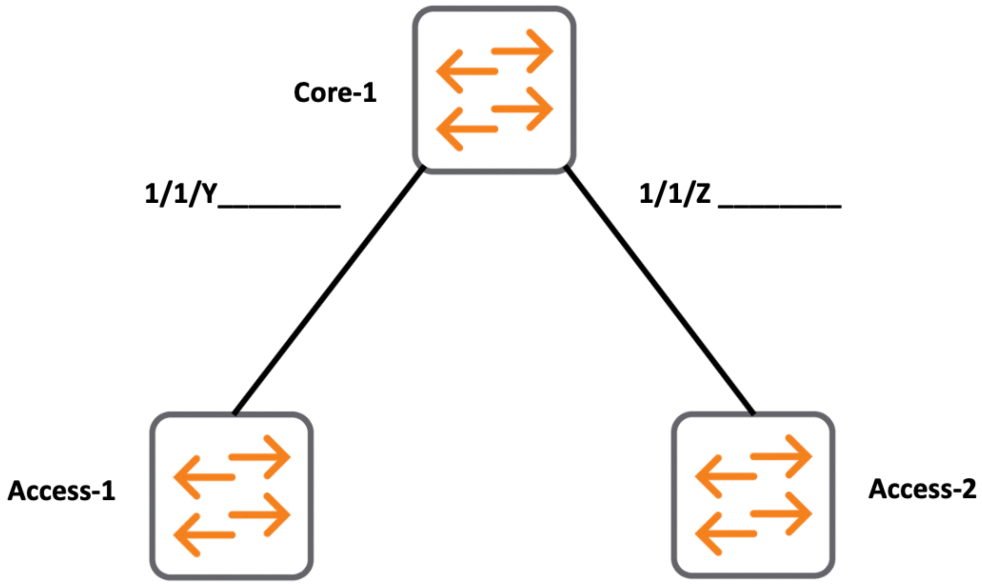

AOS-CX Switching Fundamentals
Figure 4.3-2: LLDP Discovery

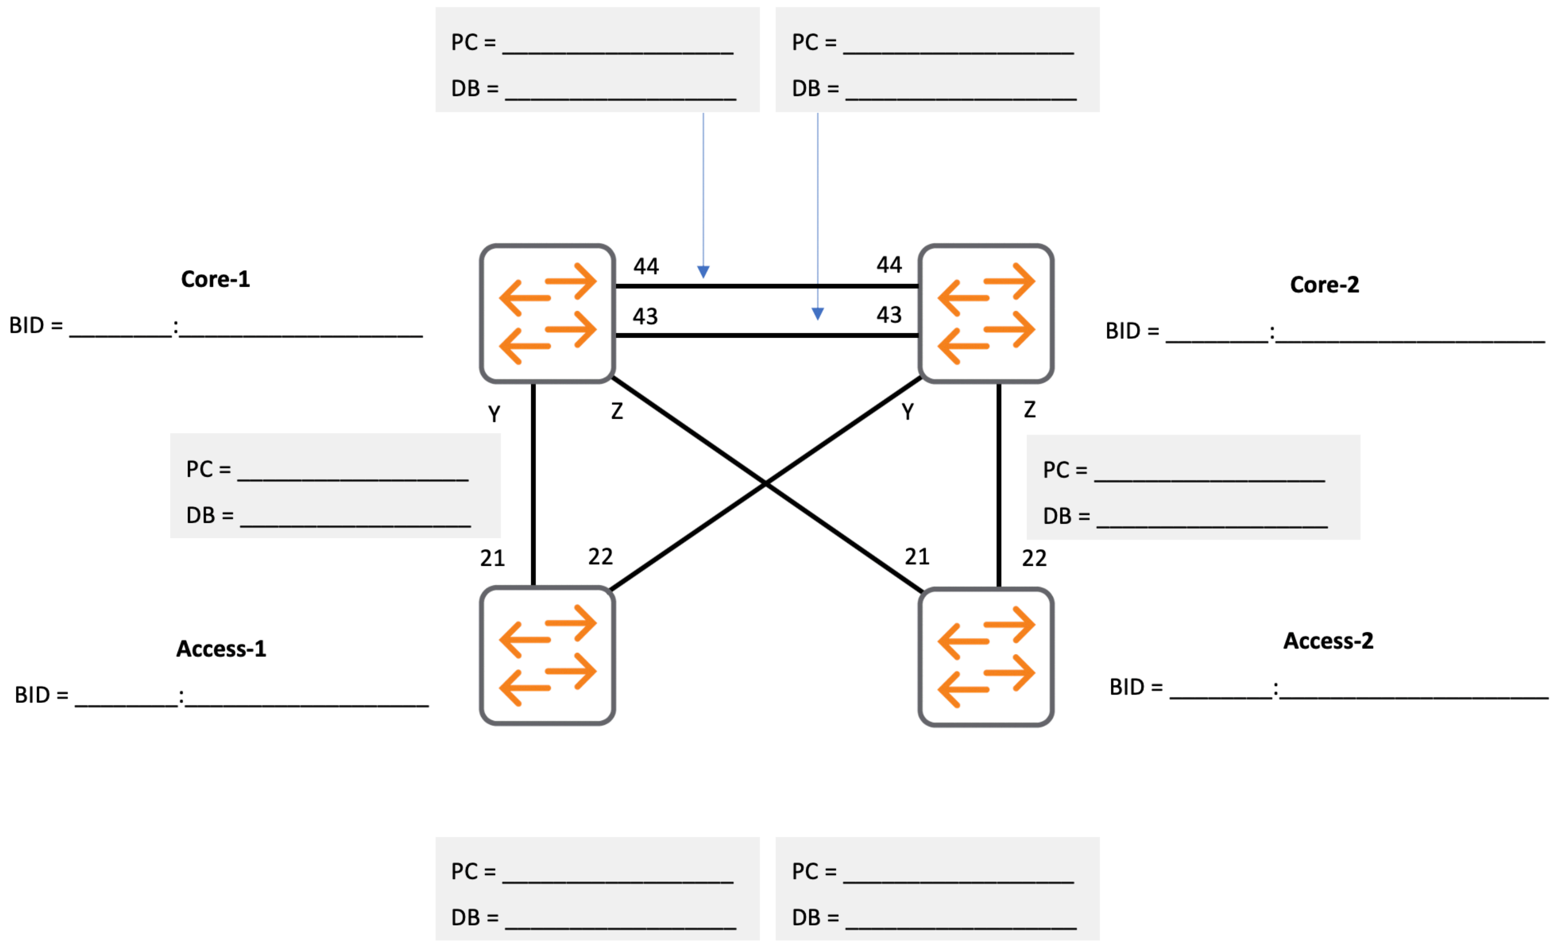

PC = ____________________
DB = ____________________

PC = ____________________
DB = ____________________

Core-1

BID = _________:_____________________

Core-2

BID = _________:_____________________

PC = ____________________
DB = ____________________

PC = ____________________
DB = ____________________

Access-1

BID = _________:_____________________

Access-2

BID = _________:_____________________

PC = ____________________
DB = ____________________

PC = ____________________
DB = ____________________

AOS-CX Switching Fundamentals
Figure 5.1-2: BIDs, Designated Bridges and costs

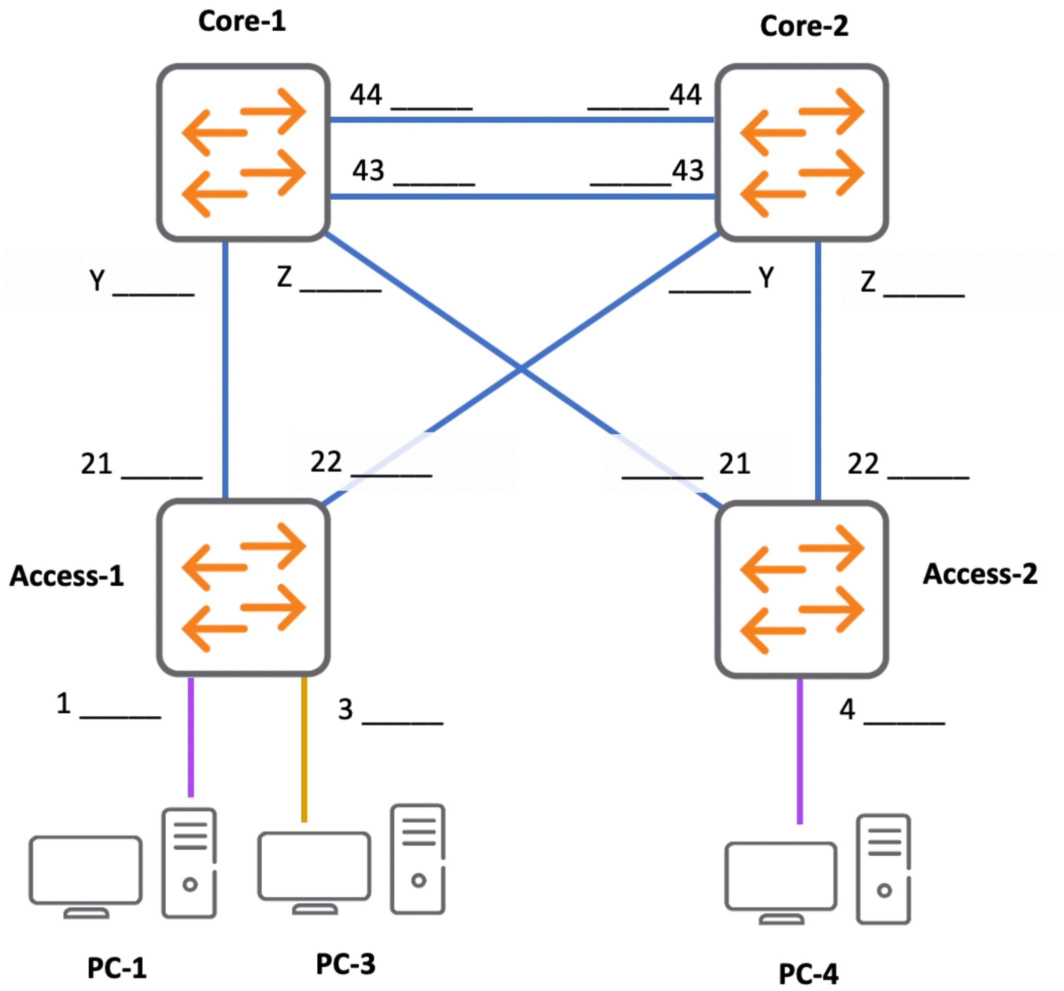

AOS-CX Switching Fundamentals
Figure 5.1-3: Devices and ports roles

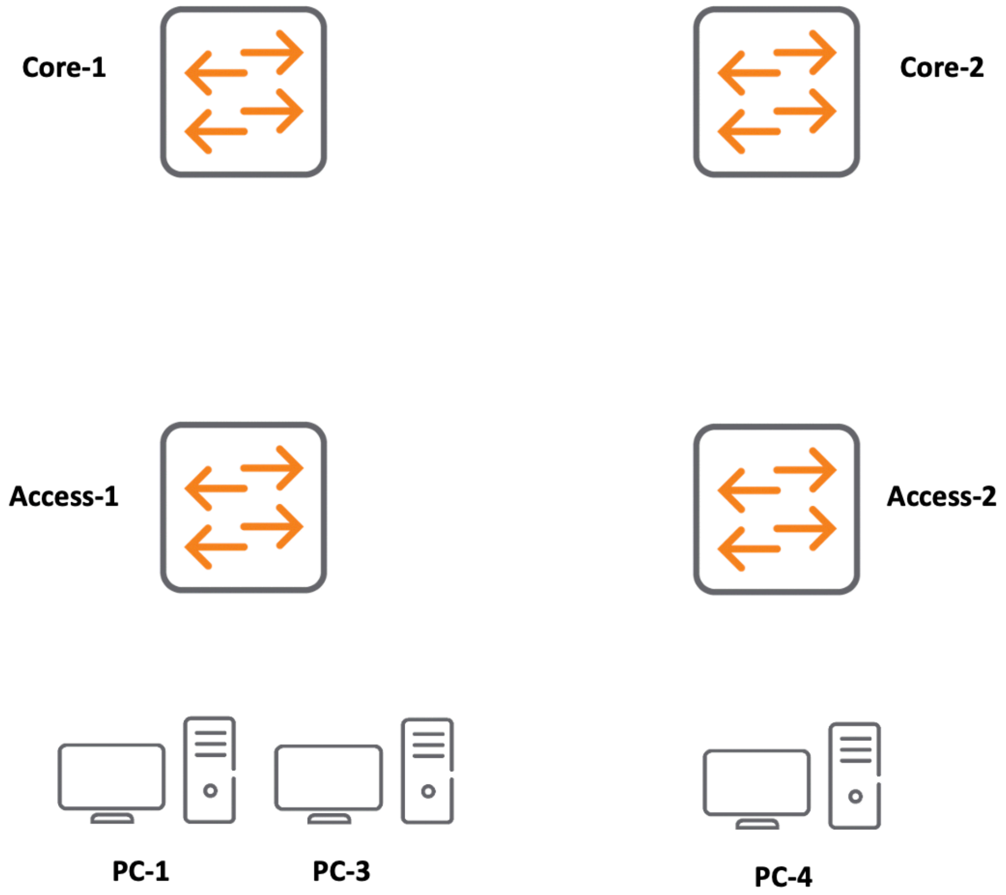

AOS-CX Switching Fundamentals
Figure 5.1-4: Drawing CST

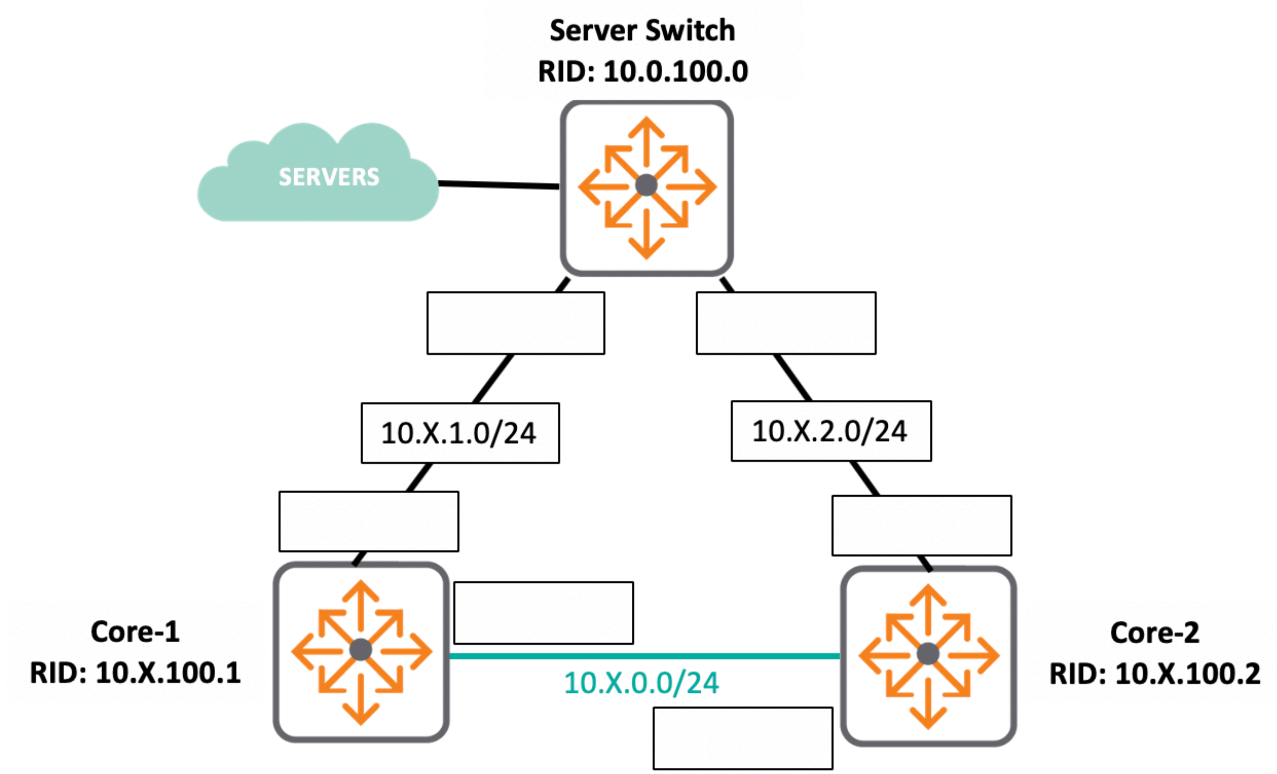

AOS-CX Switching Fundamentals
Figure 11.1-3: DRs and BDRs

6280 AMERICA CENTER DR SAN JOSE CA 95054
TEL:408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-CXF-RLABS-v20.22